IBM Cloud

IBM Cloud overview

Product guide

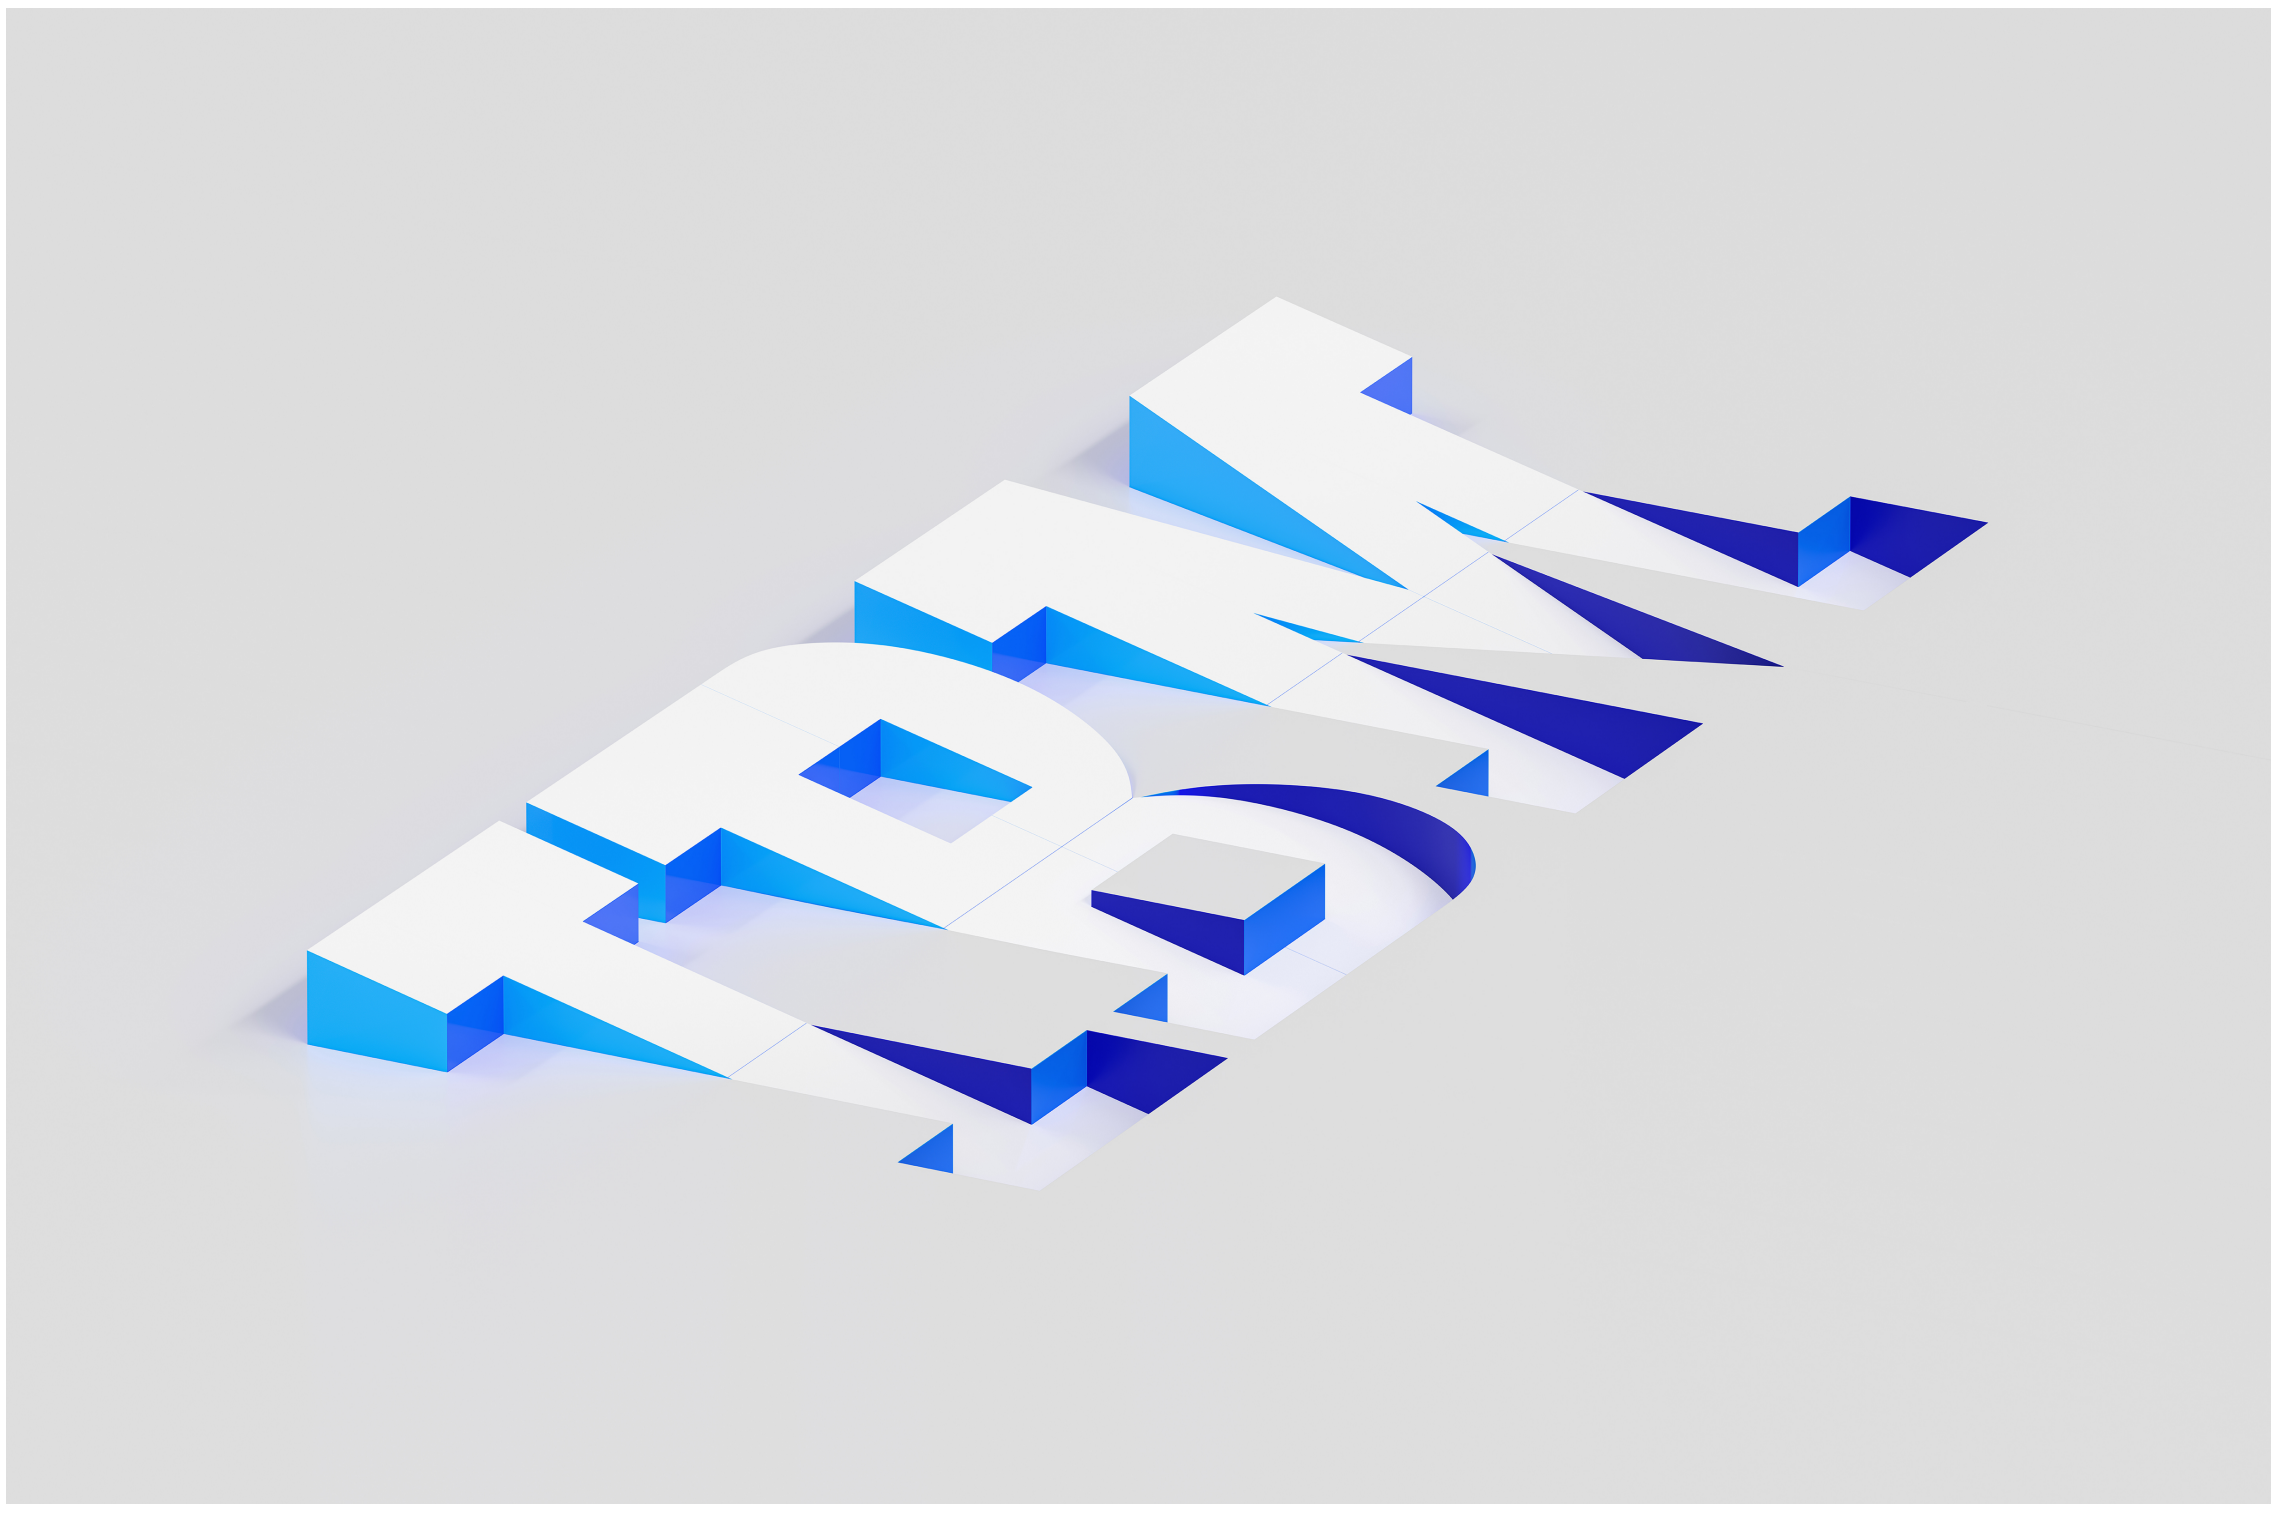

IBM

Edition notices

This PDF was created on 2025-01-22 as a supplement to *IBM Cloud overview* in the IBM Cloud docs. It might not be a complete set of information or the latest version. For the latest information, see the IBM Cloud documentation at https://cloud.ibm.com/docs/overview.

Learn about the IBM Cloud platform

What is the IBM Cloud platform?

The IBM Cloud® platform combines platform as a service (PaaS) with infrastructure as a service (IaaS) to provide an integrated experience. The platform scales and supports both small development teams and organizations, and large enterprise businesses. Globally deployed across data centers around the world, the solution you build on IBM Cloud® spins up fast and performs reliably in a tested and supported environment you can trust!

IBM Cloud provides solutions that enable higher levels of compliance, security, and management, with proven architecture patterns and methods for rapid delivery for running mission-critical workloads. Available in data centers worldwide, with multizone regions in North and South America, Europe, Asia, and Australia, you are enabled to deploy locally with global scalability.

IBM Cloud offers the most open and secure public cloud for business with a next-generation hybrid cloud platform, advanced data and AI capabilities, and deep enterprise expertise across 20 industries. Solutions are available depending on your needs for working in the public cloud, on-premises, or a combination:

- With public cloud, the resources are made available to you over the public internet. It is a multi-tenant environment, and resources like hardware and infrastructure are managed by IBM®.
- A hybrid cloud solution is a combination of public and private giving you the flexibility to move workloads between the two based on your business and technological needs. IBM uses Red Hat OpenShift on IBM Cloud, the market-leading hybrid cloud container platform for hybrid solutions that enables you to build once and deploy anywhere. With IBM Cloud Satellite, you can create a hybrid environment that brings the scalability and on-demand flexibility of public cloud services to the applications and data that runs in your secure private cloud.
- Support for multicloud and hybrid multicloud solutions is also available, which makes it easy for you to work with different vendors. IBM Cloud Paks are software products for hybrid clouds that enable you to develop apps once and deploy them anywhere.
- Virtual Private Cloud (VPC) is available as a public cloud service that lets you establish your own private cloud-like computing environment on shared public cloud infrastructure. With VPC, enterprises can define and control a virtual network that is logically isolated from all other public cloud tenants, creating a private, secure place on the public cloud.

With our open source technologies, such as Kubernetes, Red Hat OpenShift, and a full range of compute options, including virtual machines, containers, bare metal, and serverless, you have the control and flexibility that's required to support workloads in your hybrid environment. You can deploy cloud-native apps while also ensuring workload portability.

Whether you need to migrate apps to the cloud, modernize your existing apps by using cloud services, ensure data resiliency against regional failure, or use new paradigms and deployment topologies to innovate and build your cloud-native apps, the platform's open architecture is built to accommodate your use case.

What's built into the platform?

As the following diagram illustrates, the IBM Cloud platform is composed of multiple components that work together to provide a consistent and dependable cloud experience.

- A robust console that serves as the front end for creating, viewing, managing your cloud resources
- An identity and access management component that securely authenticates users for both platform services and controls access to resources consistently across IBM Cloud
- A catalog that consists of hundreds of supported products
- A search and tagging mechanism for filtering and identifying your resources
- An account and billing management system that provides exact usage for pricing plans and secure credit card fraud protection

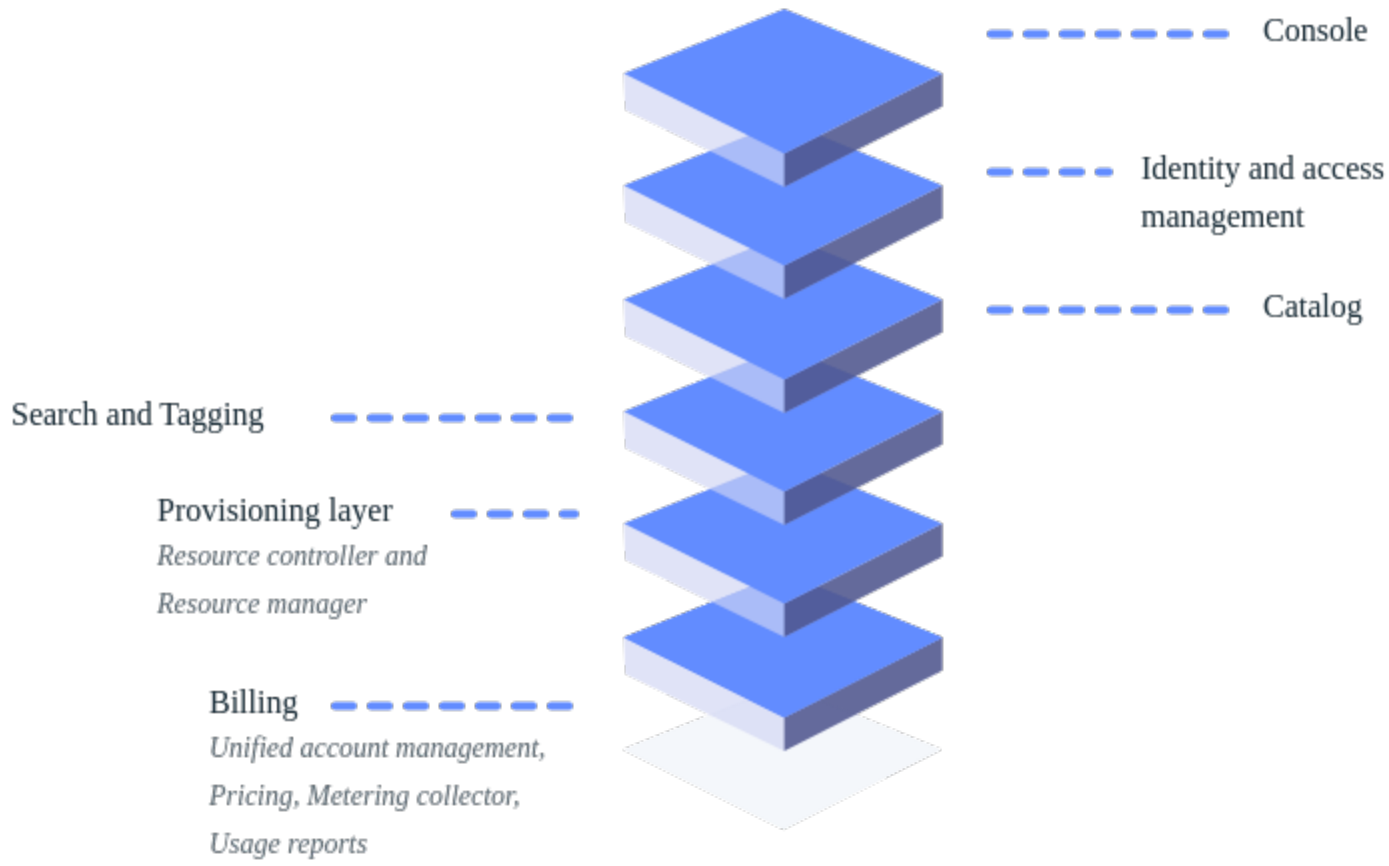

Components of the IBM Cloud platform

Whether you have existing code that you want to modernize and bring to the cloud or you're developing a brand new application, your developers can tap into the rapidly growing ecosystem of available services and runtime frameworks in IBM Cloud.

Setting up your account

If you're a developer and you're just trying out IBM Cloud, you can go straight to the catalog and browse the products that you'd like to explore. Try filtering for all Lite and Free pricing plans to test out IBM Cloud with no costs. When you're ready to get started with an environment and get apps running in production, consider setting up the basics in your account:

- Access groups for organizing users and service IDs into one entity to make assigning access a streamlined process
- Resource groups for organizing your resources to make assigning access to a set of resources quick and easy
- IAM access policies for your access groups or individual developers

For more information, see the best practices for organizing your resources and assigning access.

As a financial officer for your company, you might be interested in simplifying how you manage billing and usage across multiple teams and departments. With a Subscription account, you can create an IBM Cloud *enterprise*, which offers centralized account management, consolidated billing, and top-down usage reporting. An enterprise consists of an enterprise account, account groups, and individual accounts.

- The enterprise account is the parent account to all other accounts in the enterprise. Billing for the entire enterprise is managed at the enterprise account level.
- Account groups provide a way to organize related accounts. And, you get a unified view of resource usage costs across all accounts that are included in an account group.
- Similar to stand-alone accounts, accounts in an enterprise contain resources and resource groups and independent access permissions.

For more information, see the Enterprise account architecture white paper and the best practices for setting up an enterprise.

IBM Cloud catalog

Discover all that IBM Cloud has to offer. From services, software, and *deployable architectures* ranging from containers, compute, security, data, AI, and more, find what you need to transform your business.

The available services include options for compute, storage, networking, end-to-end developer solutions for app development, testing and deployment, security management services, traditional and open source databases, and cloud-native services. The lifecycle and operations of services are the responsibility of IBM.

You can also find a number of software products, including Cloud Paks, Terraform-based templates, Helm charts, and Operators. The preconfigured software solutions help you build faster. And, with a simplified installation process, you can get started quickly. You manage the deployment and configuration of the software on your own compute resources.

If you're looking for more robust solutions for your enterprise business goals, IBM Cloud offers deployable architectures that use cloud automation for deploying common architectural patterns that combine one or more cloud resources that are designed for easy deployment, scalability, and modularity.

And, if you're looking for help in your journey to cloud, check out our professional services. Browse your options for scheduling a consultation with technical experts depending on your needs, such as cloud migration, creating business solutions with IBM Garage, or developing a container security solution that works for you.

Tip: The catalog supports command-line interfaces (CLIs) and a RESTful API for you to use to retrieve information about existing products.

Open source offerings

In addition to the generally available offerings in the IBM Cloud catalog, you can find additional offerings for deploying and building solutions in the module registry and community registry.

IBM module registry

Discover modules that work with deployable architectures for your customization and building needs. This is a collection of assets that is separate from the IBM Cloud catalog and is governed and maintained by the process in the `terraform-ibm-modules` GitHub organization. The modules that you find here meet a specific quality checks, indicated by the Graduated or Stable badge, and are confirmed to work with deployable architectures in the IBM catalog. For more information, see Badges for modules. You can view usage information and one or more runnable examples for each module. Example's aren't deployable from the module registry, but you can copy the usage code to get started with any module.

Community registry

This is a collection of real world examples of coded industry solutions to jumpstart your building needs. The collection maintained by the owners of the originating GitHub repository and solutions might change frequently or be discontinued at short notice. Note that these solutions are't supported by the IBM Cloud Support Center. If you have questions about the solution, you can open an issue in the originating GitHub repository.

The offerings in the module registry and community registry might include the following badges:

- `TIM Approved` indicates that the solution is sourced in the `terraform-ibm-modules` GitHub organization, a collection of IBM Cloud Terraform modules.
- `IBM open source` indicates that the solution is sourced in the `ibm` GitHub organization, a collection of open source repositories provided by IBM under the open source software license.
- `Recommended` indicates that an internal team of experts has reviewed and endorses the offering for solving common or complex cloud use cases.

Searching the catalog for services

All products that are available in IBM Cloud are displayed by default in the catalog. You can filter the catalog by type to view a specific type of product, for example, only services, software, or deployable architectures. Enter keywords or set additional filters to further scope your view of the catalog. For example, if you want to deploy an analytics instance to Red Hat® OpenShift® on IBM Cloud®, you can select the **Analytics** category, and filter the results by selecting **Red Hat OpenShift** as the deployment target.

See the following table for the list of filters that you can use to search the catalog.

Option	Description
AI / Machine Learning	Products that enable systems to learn from data rather than through explicit programming
Analytics	Products that facilitate the analysis of data, typically large sets of business data, by the use of mathematics, statistics, and other means
Blockchain	Products that facilitate the process of recording transactions and tracking assets in a business network
Compute	Infrastructure resources that serve as the basis for building apps in the cloud
Containers	A standard unit of software that packages up code and all its dependencies so the app runs quickly and reliably from one computing environment to another

Databases	Products that provide some form of access to a database without the need for setting up physical hardware, installing software, or configuring for performance
Developer tools	Products that support developing, testing, and debugging software
Enterprise applications	Bundle of compatible products that deliver enterprise-grade app solutions for information sharing, automation, and agility
Integration	Products that facilitate the connection of data, apps, APIs, and devices across an organization to be more efficient, productive, and agile
Internet of Things	Products that support receiving and transferring data over wireless networks without human intervention
Logging and monitoring	Products that support storing, searching, analyzing, and monitoring log data and events. And, products that support reviewing and managing the operational workflow and processes being logged
Mobile	Products with specific or special utility for users creatings things to be used on mobile devices
Networking	Products that support or augment the linking of computers so they can operate interactively
Security	Products that provide the protection of stored data from theft, leakage, and deletion
Storage	Products that support data to be created, read, updated, and deleted

Options for filtering by category

Option	Description
EU Supported	Support for the service is restricted to IBM Cloud support team members that are located in the European Union (EU) region. This filter is available only if the EU Supported setting is enabled in the account.
Financial Services Validated	Services are designated as Financial Services Validated when the IBM Cloud service or SaaS, or independent software vendor (ISV) product, evidences compliance with the IBM Cloud Framework for Financial Services. This filter is available only for IBM products, and if the Financial Services Validated setting is enabled in the account.
HIPAA Enabled	The service is designated as HIPAA ready, meaning processing, storing, and handling Protected Health Information (PHI) in the service is supported. This filter is available only if the HIPAA Supported setting is enabled in the account.
IAM-enabled	The service is enabled to use IBM Cloud Identity and Access Management (IAM) for access control. Access policies are used to assign users and service IDs access to specific resources in an account.
Service Endpoint Supported	The service can be connected to over the IBM Cloud private network instead of the public network. Connecting directly to service endpoints doesn't require internet access, providing increased security.

Options for filtering by compliance

Option	Description
Free	The service includes monthly free allowances for only Pay-As-You-Go or Subscription accounts.
Lite	The pricing plan for the service is structured as a free quota. The quota might operate for a specific time period, for example, a month or on a one-off usage basis.

Options for filtering by pricing plan

Option	Description
Beta	The service is available for evaluation and testing purposes. Beta services aren't intended for production use.

Deprecated	Deprecated products are in the process of being withdrawn from service and are eligible to be removed after the deprecation period.

Options for filtering by release

Option	Description
IBM supported	Products that are supported by IBM Cloud.
Third party supported	Products that are provided by individual service entities.
Community supported	Products that are provided by open source communities.

Options for filtering services by support type

Option	Description
HPC	Products that enable High Performance Computing (HPC) workloads on IBM Cloud. For more information, see High-performance computing on IBM Cloud
SAP certified	An infrastructure service that is certified by SAP to run production SAP workloads. For more information, see IBM Cloud® for SAP.
Satellite enabled	A service that is enabled for use with IBM Cloud Satellite. You can run apps consistently across on-premises, edge computing, and public cloud environments. For more information, see IBM Cloud Satellite.
Quantum Technologies	A service that is compatible with quantum technologies. For more information, see IBM Quantum services.

Options for filtering by run-time environment

Tip: You can also scope your view of the catalog by using the **Provider** filter to browse by individual providers, and the **Location** filter to view products available in specific regions.

Searching the catalog for software

The following table lists the filter options you can use when searching the catalog for software.

Option	Description
AI / Machine Learning	Products that enable systems to learn from data rather than through explicit programming
Analytics	Products that facilitate the analysis of data, typically large sets of business data, by the use of mathematics, statistics, and other means
Blockchain	Products that facilitate the process of recording transactions and tracking assets in a business network
Compute	Infrastructure resources that serve as the basis for building apps in the cloud
Containers	A standard unit of software that packages up code and all its dependencies so the app runs quickly and reliably from one computing environment to another
Databases	Products that provide some form of access to a database without the need for setting up physical hardware, installing software, or configuring for performance
Developer tools	Products that support developing, testing, and debugging software
Enterprise applications	Bundle of compatible products that deliver enterprise-grade app solutions for information sharing, automation, and agility

Integration	Products that facilitate the connection of data, apps, APIs, and devices across an organization to be more efficient, productive, and agile
Internet of Things	Products that support receiving and transferring data over wireless networks without human intervention
Logging and monitoring	Products that support storing, searching, analyzing, and monitoring log data and events. And, products that support reviewing and managing the operational workflow and processes being logged
Mobile	Products with specific or special utility for users creatings things to be used on mobile devices
Networking	Products that support or augment the linking of computers so they can operate interactively
Security	Products that provide the protection of stored data from theft, leakage, and deletion
Storage	Products that support data to be created, read, updated, and deleted

Options for filtering by category

Option	Description
Cloud Paks	A cloud solution that integrates a container platform, containerized IBM middleware and open source components, and common software services for development and management.
Helm charts	A format for packaging a collection of files that describe specific configurations of infrastructure in the form of code.
OVA images	Open Virtual Appliance that contains a compressed installable version of a virtual machine.
Operators	A method of packaging and deploying a Kubernetes-native application.
Server Images	A template that is used to create instances of virtual servers.
Terraform	Infrastructure as code to deploy your application.

Options for filtering by delivery method

Option	Description
IBM Kubernetes Service	Used to create a Kubernetes cluster of compute hosts to deploy and manage containerized apps on IBM Cloud.
IBM Cloud Schematics	Used for infrastructure as code automation by using terraform templates.
Power Virtual Server	Used to create a Power server that is distinct from the IBM Cloud servers with separate networks and direct-attached storage. The internal networks are fenced but offer connectivity options to IBM Cloud infrastructure or on-premises environments.
Red Hat OpenShift	Used to create a Red Hat OpenShift on IBM Cloud cluster of compute hosts to deploy and manage containerized apps on IBM Cloud.
VMware vCenter Server	Provides deployment and management of VMware virtualized environments.
Virtual private cloud	Deploy and manage your server images on virtual private cloud as your infrastructure target.

Options for filtering by deployment target

Option	Description
Free	The service includes monthly free allowances for only Pay-As-You-Go or Subscription accounts.
Lite	The pricing plan for the service is structured as a free quota. The quota might operate for a specific time period, for example, a month or on a one-off usage basis.

Options for filtering software by pricing plan

Option	Description
IBM supported	Products that are supported by IBM Cloud.
Third party supported	Products that are provided by individual service entities.
Community supported	Products that are provided by open source communities.

Options for filtering software by support type

Option	Description
HPC	Products that enable High Performance Computing (HPC) workloads on IBM Cloud. For more information, see High-performance computing on IBM Cloud
SAP Certified	An infrastructure service that is certified by SAP to run production SAP workloads. For more information, see IBM Cloud® for SAP.
Satellite Enabled	A service that is enabled for use with IBM Cloud Satellite. You can run apps consistently across on-premises, edge computing, and public cloud environments. For more information, see IBM Cloud Satellite.
Quantum Technologies	A service that is compatible with quantum technologies. For more information, see IBM Quantum services.

Options for filtering software by run-time environment

Tip: You can also scope your view of the catalog by using the **Provider** filter to browse by individual providers, and the **Location** filter to view products available in specific regions.

Searching the catalog for deployable architectures

You can search our growing catalog of deployable architectures to find preassmbled cloud automation solutions that solve common enterprise business needs, for example a secure infrastructure layer for highly regulated industries, such as financial services.

The following table lists the filter options that you can use when searching the catalog for deployable architectures.

Option	Description
AI / Machine Learning	Products that enable systems to learn from data rather than through explicit programming
Analytics	Products that facilitate the analysis of data, typically large sets of business data, by the use of mathematics, statistics, and other means
Blockchain	Products that facilitate the process of recording transactions and tracking assets in a business network
Compute	Infrastructure resources that serve as the basis for building apps in the cloud
Containers	A standard unit of software that packages up code and all its dependencies so the app runs quickly and reliably from one computing environment to another

Databases	Products that provide some form of access to a database without the need for setting up physical hardware, installing software, or configuring for performance
Developer tools	Products that support developing, testing, and debugging software
Enterprise applications	Bundle of compatible products that deliver enterprise-grade app solutions for information sharing, automation, and agility
Integration	Products that facilitate the connection of data, apps, APIs, and devices across an organization to be more efficient, productive, and agile
Internet of Things	Products that support receiving and transferring data over wireless networks without human intervention
Logging and monitoring	Products that support storing, searching, analyzing, and monitoring log data and events. And, products that support reviewing and managing the operational workflow and processes being logged
Mobile	Products with specific or special utility for users creatings things to be used on mobile devices
Networking	Products that support or augment the linking of computers so they can operate interactively
Security	Products that provide the protection of stored data from theft, leakage, and deletion
Storage	Products that support data to be created, read, updated, and deleted

Options for filtering by category

Tip: You can also scope your view of the catalog by using the **Provider** filter to browse by individual providers and the **Industry** filter to view products catered for certain industries.

Pricing and billing

You can view the pricing details for each service when you're browsing the catalog. If you choose a service plan with a paid plan, you can estimate your costs by using the cost estimator tool. For more information, see Estimating your costs.

IBM Cloud billing provides multiple services that ensure the IBM Cloud platform can securely manage pricing, accounts, usage, and more.

Account management

Account management maintains the billing relationship with the customer. Each account is a billing entity that represents a customer. This service controls account lifecycle, subscription, user relationship, and organization.

Usage metering

With usage metering, service providers can submit metrics that are collected for resource instances that are created by IBM Cloud users. Third-party service providers that deliver an integrated billing service are required to submit usage for all active service instances every hour.

Usage reports

Usage reports return the summary for the account for the specified month. Account billing managers are authorized to access the reports.

Managing security and compliance

The IBM Cloud® Security and Compliance Center offers a single location where you can validate that your resources are meeting continuous security and compliance.

You can create profiles and config rules to ensure that specific areas of your business adhere to your defined requirements or industry regulations. From the Security and Compliance Center dashboard, you can download detailed reports that you can use to provide evidence to stakeholders or external auditors. The Security and Compliance Center also offers security insights that you can use to detect potential threats when observing your account activity. For more information, see Getting started with Security and Compliance Center.

Creating resources

The resource controller is the next-generation IBM Cloud platform provisioning layer that manages the lifecycle of IBM Cloud resources in your account. Resources are created globally in an account scope. The resource controller supports the creation of resources both synchronously and asynchronously. Examples of resources include databases, accounts, processors, memory, and storage limits.

In general, resources that are tracked by the provisioning layer are intended to associate usage metrics and billing, but that isn't always the case. In some cases, the resource might be associated with the provisioning layer to ensure that its lifecycle can be managed along with the account lifecycle. The resource controller uses IBM Cloud Identity and Access Management (IAM) for authentication and authorization of actions that are taken against the provisioning layer.

The resource controller provides common APIs to control the lifecycle of resources from creating an instance to creating access credentials to removing access to deleting an instance.

Managing your resources

A collection of resources is managed by resource groups. A resource group is associated with your account. All IBM Cloud resources must be assigned to a resource group. When you create an account, a default resource group is created for you. All IBM Cloud IAM-enabled resources must be created within a resource group. If you have a Lite account, you can have only one resource group, but with a a Pay-As-You-Go or Subscription account, you can create more than one resource group. If an account is suspended, the corresponding resource group is suspended as well, and all resources within the resource group are suspended.

Managing Infrastructure as Code (IaC) deployments with projects

IBM Cloud *projects* are a named collection of configurations that are used to manage related resources and Infrastructure as Code (IaC) deployments across accounts. They enable teams to configure, deploy, and monitor deployments by using DevOps best practices. If you select a deployable architecture from the catalog, you can add it to a project to configure and deploy it into your different environments. For more information, see Learn about IaC deployments with projects.

Searching and tagging resources

The search service is a global and shared resource properties repository that is integrated within the IBM Cloud platform. It is used for storing and searching a cloud resource's attributes, and it categorizes and classifies resources. Resources are uniquely identified by a Cloud Resource Name (CRN) identifier. The properties of a resource include tags and system properties. Both properties are defined within an IBM Cloud billing account, and span across many regions.

This service also manages tags that are associated with a resource. You can create, delete, search, attach, or detach tags with the Tagging API. Tags are uniquely identified by a CRN identifier. Tags have a name, which must be unique within a billing account. You can create tags in key:value pairs or label format.

Monitoring your resources

Observability offers a single location where you can monitor and observe your applications and services in IBM Cloud.

With the IBM® Log Analysis service, you can add log management capabilities to your IBM Cloud architecture and you can manage system and application logs. It offers advanced features to monitor and troubleshoot, define alerts, and design custom dashboards. For more information, see Getting started with IBM Log Analysis.

You can gain operational visibility into the performance and health of your applications, services, and platforms with the IBM Cloud Monitoring service. It offers a full stack telemetry with advanced features to monitor and troubleshoot, define alerts, and design custom dashboards. For more information, see Getting started with IBM Cloud Monitoring.

Monitoring your account

Use the IBM Cloud® Activity Tracker service to monitor the activity of your IBM Cloud account, investigate abnormal activity and critical actions, and comply with regulatory audit requirements. In addition, you can be alerted on actions as they happen. The events that are collected comply with the Cloud Auditing Data Federation (CADF) standard. For more information, see Getting started with IBM Cloud Activity Tracker.

Viewing status

The IBM Cloud Status page is the central place to find all unplanned incidents, planned maintenance, announcements, and security bulletin notifications about key events that affect the IBM Cloud platform. You can filter these categories by selecting specific locations, components, types of ongoing events, or by using keyword searches. For more information, see Viewing cloud status.

Notification preferences

Depending on your IBM Cloud account type, you can choose to receive email notifications about IBM Cloud platform-related items and resource-related items from the Notification preferences page. Platform-related items include announcements, billing and usage, and ordering. Resource-related items include incidents, maintenance, security bulletins, and resource activity. For more information, see Setting email preferences for notifications.

Service availability in the IBM Cloud console, APIs, SDKs, CLIs, and Terraform

IBM Cloud® is composed of multiple services that are available in various interfaces for a consistent cloud experience.

Services that are available in the IBM Cloud console are also available through APIs, SDKs, CLIs, and Terraform at their *general availability (GA)* release. SDKs are, at a minimum, available in the Java, Node, Python, and Go programming languages.

- IBM Cloud catalog
- IBM Cloud API and SDK reference library
- IBM Cloud CLI reference
- IBM Cloud Terraform reference

For every service, release notes and change logs are recorded for each update that is delivered. As of 1 January 2022, services will maintain release notes for a minimum of three years. You can view release notes in each service's documentation, or you can go to the IBM Cloud Status UI to view release notes for all services in one place where you can search, filter by component, or filter by date.

What are the IBM Cloud prerequisites?

The prerequisites for using the IBM Cloud® platform are limited, but we do have a few.

Browsers

The following table specifies the minimum required browser software for IBM Cloud.

Browser	Minimum Required Browser Software
Chrome	Latest version -1 for your operating system
Firefox	Latest regular -1 and ESR versions for your operating system, see Mozilla Firefox Extended Support Release for more details
Edge	Latest version -1 for Windows
Safari	Latest version -1 for Mac

Minimum required browser software for IBM Cloud

Command-line interface

The IBM Cloud command-line interface (CLI) is constantly changing. To see the latest CLI version, go to Getting started with the IBM Cloud CLI.

Getting started on IBM Cloud checklist

Use these onboarding checklists to create and configure your IBM Cloud® account. This guide is intended to help you quickly navigate the available documentation to get your account set up, secure your cloud resources, track costs and billing, set up on-premises cloud connectivity, and help you efficiently meet your business needs in your IBM Cloud account.

Tip: Download the single page IBM Cloud Quick Start Guide to share a consolidated version of this checklist within your organization to help accelerate migrating workloads to IBM Cloud.

Explore the platform

If you have little to no experience with IBM Cloud, or you need a refresher on the IBM Cloud platform, start with the following tasks.

Task	Description
• ☐ **Learn about the IBM Cloud platform**	Review the What is the IBM Cloud platform? documentation to get familiar.
• ☐ **Understand your options for working with resources in IBM Cloud**	To create and manage resources in IBM Cloud, you can use any of the following tools: • IBM Cloud console • IBM Cloud Command Line Interface (CLI) • IBM Cloud Shell. • IBM Cloud Schematics • Each service offers a set of APIs and SDKs available as part of their documentation.
• ☐ **Learn how to navigate the IBM Cloud console**	The console is the user interface that you use to create and manage all your IBM Cloud resources. You can create a free account, log in, access documentation, access the catalog, view pricing information, get support, or check the status of IBM Cloud components. For more information, see Navigating the IBM Cloud console.
• ☐ **Get started with the IBM Cloud CLI**	The command line interface (CLI) is a set of plug-ins and tools that you can use to create and manage your resources. For more information, see Getting started with the IBM Cloud CLI.
• ☐ **Get started with IBM Cloud Shell**	The Cloud Shell gives you a personal cloud-based shell workspace with the full IBM Cloud CLI and more command-line tools with no installation needed. Learn more about Working in IBM Cloud Shell.
• ☐ **Get started with IBM Cloud Schematics**	Schematics provides powerful set of Infrastructure as Code (IaC) tools - Terraform, Ansible, Helm - as a service to program your cloud infrastructure. Schematics can run your end-to-end automation to build one or more stacks of cloud resources, manage their lifecycle, manage changes in their configurations, deploy your app workloads, and perform Day 2 operations. Review the Getting started: IBM Cloud Schematics documentation.

• ☐ **Get started with APIs**	IBM Cloud services provide APIs that comply with OAuth 2.0 authentication standards and accept bearer tokens that are provided by IBM Cloud's Identity and Access Management (IAM) service. Explore the API docs for the services that you plan on using.

Getting started tasks for exploring the platform

Set up accounts and enterprises

This checklist is for administrators who are responsible for creating and setting up an account structure in IBM Cloud and enabling users within their company to create and manage cloud resources. IBM Cloud offers you the ability to create a stand-alone account and an enterprise.

Stand-alone account

This type of account allows an account owner, for example a department or business unit administrator to add users to the account, assign access roles and permissions, manage billing and payments, and more.

Enterprise

An enterprise manages the billing for the entire company, with usage costs from multiple accounts being rolled up and paid for from the enterprise account. Accounts that are created as part of an enterprise are just like stand-alone accounts, but the main difference is that these accounts don't manage their own billing or payments.

Use the following checklist to track all of the tasks to create and configure your IBM Cloud account or enterprise.

Task	Description
• ☐ **What's in an account?**	Your IBM Cloud account includes many interacting components and systems for resource, user, and access management. Understanding concepts like how certain components are connected or how access works help you effectively set up your account. For more information, see What's in an account.
• ☐ **Determine if you need an enterprise**	When you create an enterprise to manage your billing, you can move existing stand-alone accounts to it, or create new accounts as needed. Consider the following when determining whether you need an enterprise or a stand-alone account: • In an enterprise, subscription discounts and cloud credits are available to all accounts that are in the enterprise. • Stand-alone accounts control their own billing. If your company is globally distributed, you might have a mix of multiple enterprises and stand-alone accounts to support regional billing requirements. Review the What is an enterprise? documentation to determine whether you need an enterprise.
• ☐ **Create and configure your IBM Cloud account**	Even if you plan on using an enterprise, you need to create an IBM Cloud account. You can create your account by going to the account registration page and providing an email address and other additional information. The email address that is used to register becomes the account owner, but you can change this if required later on by following the steps in Transferring ownership of your account. If you have a Google account, you can use your Google credentials to create an IBM Cloud account from the account registration page. For more information, see Using a Google ID. When setting up an account for your company or organization, it is best to use a functional ID, some teams call them service accounts, associated with your company. Keep in mind that you will need to monitor for automated emails sent to this email address for warnings about service usage, services being deprecated, new services available, and more. When you log in to the account for the first time, you are required to provide a credit card or subscription code to complete your account set up. Later, you can add users by inviting them to the account or by federating to your own corporate directory. Users that are added to your account are not required to create their own account. For more account-related FAQs, visit FAQ library for managing your account, resources, and access.

• ☐ **Set up SAML federation**	By default, when you create an account you use an IBMid for user identity. IBMid is the ID as a Service (IDaaS) from IBM® used to access IBM web-based services, including IBM Cloud resources. The IBMid is based on your company's email address and a password that is managed by IBMid. IBMid allows you to federate to your own corporate directory or a third-party Identity Provider (IdP) service that you might already be using such as Okta. Federating to your own directory simplifies the process of adding users to your account as they will not require an IBMid with a separate password. However, there are cases where it might not be feasible for you to use IBMid federated to your corporate directory. An alternative is to create an IBM Cloud App ID instance in your account and connect it to your chosen identity provider. Consider the following options that are available to you: • **No federation**: Your company email domain is not federated with IBMid, and you choose to keep it as is. Any user that uses your email domain can create their own IBMid, and the password they create is managed by IBMid. They can invite other users with an IBMid to their IBM Cloud account. • **Federate with IBMid**: If your company's email domain is already federated, you can start configuring access to your account. If it is not already federated, a manual process with the IBMid federation team is required to establish federation. The decision to federate to your company's own directory needs to involve the person in your company that can make company-wide decisions when it comes to connecting to external parties for identity services. You need to ensure that this person is included in the process. Federating with IBMid can have an impact on web services that your company already uses with IBM. • **Federate with App ID and your IdP** : This option requires the creation of an App ID instance. It is a self-service option with a low-usage fee. Choosing this option requires a custom URL to log in to IBM Cloud. In addition, an App ID instance and configuration of the federation to your IdP is required for every IBM Cloud account. For a deeper dive, use the following guides and documentation: • IBM Cloud SAML Federation Guide • IBM Cloud login sequences. Use the following documentation topics when you are ready to federate: • IBMid Enterprise Federation Adoption Guide • Enable authentication from an external identity provider using App ID.
• ☐ **Activate a subscription code**	With an IBM Cloud subscription, you get discounted usage for platform services and support by committing to a minimum spending commitment for a certain period of time. After you buy a subscription for platform or support credit, you must add the credit to your account by applying a subscription code to your stand-alone account or enterprise. Applying the code ensures that the credit is added to your account, and you don't have unexpected overage charges. Make sure to add any purchased subscriptions to your account before creating resources. For more information, see applying subscription codes.
• ☐ **Create and configure an enterprise**	When you create an enterprise, the account that you used to initiate the process is automatically added to the enterprise, and a new enterprise account is created to manage the billing for the enterprise. Follow the steps in the Setting up an enterprise documentation to create an enterprise. Keep in mind that when using an enterprise, users within each account in the enterprise can create, use, and collaborate on resources just as they can in a stand-alone account.
• ☐ **Assign enterprise access to users**	If you chose to create an enterprise, you might require that administrators manage the enterprise performing functions, such as creating account groups, creating and managing accounts. Review the access that is required, and add users as needed. For more information, see Assigning access for enterprise management.
• ☐ **Create resource groups**	Use resource groups to organize an account's resources for access control and billing purposes. For example, creating a resource group per project allows costs to be tracked at the project level even when your resources are distributed across regions. Learn more from the What makes a good resource group strategy? best practices guide, and when you are ready, create your resource groups.

• ☐ **Set up email notification preferences**	Depending on your account type, you can choose to receive email notifications about IBM Cloud platform-related items such as announcements, billing and usage, additional notification preferences and ordering. You can update your preferences to receive email notifications about resource related, these notifications are for only the resources you use. For more. For more information, see Setting email preferences for notifications.

Secure your account and resources

As an account owner or a user in the account with the required access roles that are assigned for managing account settings and IAM access for users, you can complete tasks on the following checklist. Review and complete the following tasks to learn about how you can ensure the security of your account and resources.

Task	Description
• ☐ **Set up multifactor authentication**	Multifactor authentication (MFA) adds an extra layer of security to your account by requiring all users to authenticate by using an extra authentication factor beyond an ID and password. This is also commonly known as two-factor authentication (2FA). Review the types of multifactor authentication that can be enabled for your account. Follow the steps in enabling MFA for your account to configure the setting that is most appropriate for your company.
• ☐ **Identity and Access Management (IAM) overview**	Learn about what IBM Cloud IAM is, how IAM works, what features are available, and how to access the console, CLI, and APIs to work with IAM in your account. Learn more about how IBM Cloud IAM works.
• ☐ **Track your account activities by setting up an IBM Cloud Activity Tracker instance**	You can use the IBM Cloud Activity Tracker service to track how users and applications interact with IBM Cloud resources that you have created. Follow the instructions in Provisioning an Activity Tracker instance to create an instance in your wanted cloud region. For environments seeking to maintain **Financial Services Validated** status on IBM Cloud, you should instead follow the instructions in Use Activity Tracker Event Routing. You must create an instance of the Activity Tracker service in the Frankfurt (eu-de) region to start tracking IAM events. Learn more about the auditing events that are sent to that region.
• ☐ **Configure streaming of IBM Cloud Activity Tracker (optional)**	You can stream data from an IBM Cloud Activity Tracker instance to another IBM Cloud Activity Tracker instance across regions or to other corporate tools such as Security Information and Event Management (SIEM) tools. Learn more about streaming data.
• ☐ **Secure your resources with context-based restrictions**	Context-based restrictions give account owners and administrators the ability to define and enforce access restrictions for IBM Cloud resources based on the network location of access requests. These restrictions work with traditional IAM policies, which are based on identity, to provide an extra layer of protection. Since both IAM access and context-based restrictions enforce access, context-based restrictions offer protection even in the face of compromised or mismanaged credentials. Learn more about what context-based restrictions are and follow the guide to secure your resources using context-based restrictions.

• ☐ **Encrypt and protect your data**	You can choose from various secrets management and data protection products that help you to protect your sensitive data and centralize your secrets. Review Which data protection service is best for me? to better understand the different offerings that you can use with IBM Cloud to protect your application secrets. Use the following guides to create and configure your Secrets Manager instance: • Store your secrets in a Secrets Manager instance. • Encrypt your data with Key Protect. • Encrypt your data with Hyper Protect Crypto Services.
• ☐ **Security and compliance**	IBM Cloud provides a secure cloud platform that you can trust. IBM Cloud compliance results from a platform and services that are built on best-in-industry security standards, including GDPR, HIPAA, ISO 9001, ISO 27001, ISO 27017, ISO 27018, PCI, SOC2, and others. For more information, see Understanding compliance in IBM Cloud.

Getting started tasks for securing your account and resources

Manage billing and usage

Account owners and users with the administrator role on the billing account management service have access to monitor and manage billing, usage, invoices, payments, and more. Complete the following checklist to get familiar with the best practices and tools that you use to manage and track billing and usage in the account.

Task	Description
• ☐ **How to manage billing and usage**	Learn about the IBM Cloud billing options and tools that you can use to track your usage and manage invoicing and payments. Check out the How can I manage billing and usage in IBM Cloud video.
• ☐ **Understand suspended billing on Virtual Servers for VPC**	Suspended billing is an option that is available for virtual server instances running on VPC. It is not available for bare metal or dedicated hosts on VPC. With suspended billing, there are some resources, for example network and storage that continue billing. Learn more about suspended billing for VPC.
• ☐ **Configure spending alerts for your account**	You can enable spending notifications for Pay-As-You-Go or Subscription accounts, these alerts are configurable for the entire account or for individual services. Set up spending notifications.
• ☐ **Manage subscriptions and usage**	Setting up a subscription is covered in Set up accounts and an enterprise checklist. You can add more subscriptions and monitor subscription usage on the Commitments and subscriptions page in the IBM Cloud console. Learn more about managing subscriptions.
• ☐ **Manage commitments, subscriptions, and usage**	Setting up a commitment or subscription is covered in Set up accounts and an enterprise checklist. You can add more commitments or subscriptions as well as monitor usage on the Commitments and Subscriptions page in the IBM Cloud console. Learn more about Enterprise Savings Plan and Managing subscriptions.

• ☐ **View invoices and build your own reports**	To manage and view your invoices, visit the Invoices page from the billing and usage dashboard in the IBM Cloud console. See Viewing your invoices. You can also build your own reports by using the API and SDK that are available. • Usage Reports API/SDK • Enterprise Billing Units API/SDK. • Enterprise Usage Reports API/SDK.

Getting started tasks for managing billing and usage

Connect your network to IBM Cloud

As the need for global reach and 24/7 operations of web applications increases, the need to host services in multiple cloud data centers increases too. Data centers across multiple locations provide resilience in the case of a geographic failure and bring workloads closer to globally distributed users, which reduces latency and increases perceived performance. The IBM Cloud network enables users to link workloads hosted in secure private networks across data centers and locations. Use the following checklist to review the available options and to connect your existing on-premises environments to IBM Cloud.

Task	Description
• ☐ **Enable virtual routing and forwarding (VRF)**	Dispersed cloud resources are resources in more than one location or in more than one subnet or VLAN. These types of resources require a routing function to communicate among themselves, even within a private network context. New accounts that are created in IBM Cloud have a "multiple isolation" tenancy communication option, which is often called a customer VRF, enabled. Verify that it is enabled in your account, or enable it after confirming potential service disruptions if you have existing resources in your account. Follow the steps in the Enabling VRF in the console documentation.
• ☐ **Enable service endpoints**	It is also recommended to enable your account for using service endpoints. When IBM Cloud service endpoints are enabled in your account, you can choose to expose a private network endpoint when you create a resource. You can then connect directly to this endpoint over the IBM Cloud private network rather than the public network. Because resources that use private network endpoints don't have an internet-routable IP address, connections to these resources are more secure. Follow the steps that are outlined in Enabling service endpoints.
• ☐ **Determine the interconnectivity solution for you**	Virtual Private Networking (VPN) access enables you to manage all servers and services that are associated with your account, remotely, over the IBM Cloud private network. Depending on the infrastructure that you plan to use in IBM Cloud, there are a few options available for VPN: • VPN for classic infrastructure • Client-to-site VPN for VPC infrastructure • Site-to-site VPN for VPC infrastructure Direct Link is an alternative to a traditional site-to-site VPN solution. It can provide higher-throughput connectivity between a remote network and IBM Cloud environments. Use this decision tree to help you decide which Direct Link solution works best for you. For more information, see How do I know which Direct Link solution to order?
• ☐ **About IBM Cloud Direct Link (2.0)**	IBM Cloud Direct Link offerings provide connectivity from an external source into a customer's IBM Cloud private network. Direct Link can be viewed as an alternative to a traditional site-to-site VPN solution, which is designed for customers that need more consistent, higher-throughput connectivity between a remote network and their IBM Cloud environments. When selecting to use Direct Link, it is recommended that most customers choose Direct Link 2.0. Learn more about Direct Link (2.0).
• ☐ **Order IBM Cloud Direct Link Dedicated**	Allows customers to terminate a single-tenant, fiber-based cross-connect into the IBM Cloud network. This offering can be used by customers with colocation premises that are next to IBM Cloud PoPs and data centers, as well as network service providers that deliver circuits to customer on-premises or other data centers. For more information, see Ordering IBM Cloud Direct Link Dedicated.
• ☐ **Order IBM Cloud Direct Link Connect**	Offers private access to your IBM Cloud infrastructure and to any other clouds linked to your service provider through your local IBM Cloud data center. This option is perfect for creating multi-cloud connectivity in a single environment. IBM connects customers to the IBM Cloud private network, by using a shared bandwidth topology. As with all Direct Link products, you can add global routing that enables private network traffic to all IBM Cloud locations. For more information, see Ordering IBM Cloud Direct Link Connect.

• ☐ **Identify potential conflicts with the IBM network reserved IPs (applicable to Direct Link 1.0 only)**	When selecting to use Direct Link, it is recommended that most customers choose Direct Link 2.0, since it provides more options and flexibility. If you have a real need for Direct Link 1.0, make sure to review the list of IBM network reserved IPs that might conflict with your on-premises environment. Review IBM Cloud IP ranges.
• ☐ **Coordinate with IBM Cloud Design Engineering when you have special requirements**	If you are planning on using any of the following: GRE or IPsec tunnels, BCR pairing, multi VLAN tenants, custom inbound or outbound ACLs, ASN pre-pend, and static routes, open a support case and request assistance from IBM Cloud Design Engineering to confirm that your requirements are in line with your Direct Link selections.

Enable logging and monitoring

Analyze logs, collect metrics, and configure near real-time alerts on your cloud resources and applications.

Task	Description
• ☐ **Create platform logs**	You can use IBM Cloud Log Analysis to manage operating system logs, application logs, and platform logs in the IBM Cloud. Platform logs are logs that are exposed by enabled services and the platform in IBM Cloud. You must configure a logging instance in a region to monitor these logs. Learn more about configuring IBM Cloud Log Analysis platform logs.
• ☐ **Create a dedicated logging instance (optional)**	IBM Cloud Log Analysis offers administrators, DevOps teams, and developers advanced features to filter, search, and tail log data, define alerts, and design custom views to monitor application and system logs. Learn more about getting started with IBM Cloud Log Analysis.
• ☐ **Configure streaming of IBM Cloud Log Analysis**	You can stream data from an IBM Cloud Log Analysis instance to another IBM Cloud Log Analysis instance across regions or to other corporate tools such as Security Information and Event Management (SIEM) tools. Learn more about streaming data.
• ☐ **Configure logging agent**	The logging agent collects and forwards logs to your IBM Cloud Log Analysis instance. After you provision an IBM Cloud Log Analysis instance, you must configure a logging agent for each log source, for example in classic infrastructure or VPC infrastructure, that you want to monitor. Learn more about logging agents.

• ☐ **Configure Monitoring**	IBM Cloud Monitoring is a cloud-native, and container-intelligence management system that you can include as part of your IBM Cloud architecture. Use it to gain operational visibility into the performance and health of your applications, services, and platforms. It offers administrators, DevOps teams, and developers full stack telemetry with advanced features to monitor and troubleshoot, define alerts, and design custom dashboards. Learn more about getting started with Monitoring.
• ☐ **Configure monitoring agent**	Depending on the compute resource type that you are monitoring, follow these guides to configure the monitoring agent in classic infrastructure or VPC infrastructure: • Working with the Kubernetes agent • Working with the Red Hat OpenShift on IBM Cloud agent • Working with the Linux agent. • Working with the Docker agent.
• ☐ **Configure streaming of IBM Cloud Monitoring**	You can use IBM Cloud Monitoring to push a set of selected metrics to a Kafka service such as IBM® Event Streams for IBM Cloud®. For more information, see Streaming metrics to a Kafka service.

Streamline access management with identities, groups, and policies

Use the following checklist to create user and service identities in your accounts. Then, create access and resource groups for organizing users and resources to streamline the access management process. Account owners or users with the administrator role on all IAM account management services, which includes services such as the user management, access groups, identity service, and more can complete these tasks.

Task	Description
• ☐ **Create access groups**	Access groups are used to organize a set of users and service IDs into a single entity to enable the assignment of policies to the group instead of assigning the same access multiple times per individual user or service ID. A logical way to assign access to resource groups is by creating one access group per required level of access. Then, each access group is mapped to the needed resource groups. Examples of access groups that you should consider would be Admin and Developer. Check out the What makes a good access group strategy best practices documentation.
• ☐ **Configure access groups for Activity Tracker and Log Analysis**	You can configure, control, and manage events and logging data that is available to users by configuring groups in the logging instance. Access groups provide extra security by allowing users to see only a subset of auditing events, as opposed to all auditing events that are generated in the account. For example, you might grant a group of users access to see auditing events that are only related to development services in the account. Learn more by reviewing Using groups to control data access and RBAC, groups, and IAM integration.
• ☐ **Organize resource groups**	A resource group is a way for you to organize your account resources in customizable groupings so that you can quickly assign users access to multiple resources at a time. It is generally recommended to create a resource group per project. Learn more about managing resource groups.

• ☐ **Invite users to your account**	Users are invited to an account and given access to the resources. Use IAM to invite users, cancel invitations, or resend a pending invitation. You can invite a single user or multiple users. Start inviting users to an account.
• ☐ **Set up trusted profiles**	Trusted profiles are used to automatically grant federated users access to your account with conditions based on SAML attributes from your corporate directory. Trusted profiles can also be used to set up fine-grained authorization for applications that are running in compute resources. This way, you aren't required to create service IDs or API keys for the compute resources. Learn more about creating trusted profiles.
• ☐ **Create service IDs**	A service ID identifies a service or application similar to how a user ID identifies a user. You can create a service ID and use it to enable an application outside of IBM Cloud access to resources in your account. Learn more about creating and working with service IDs. If you have configured a Secrets Manager instance as described in Encrypt and protect your data, you can use that instance to dynamically generate a service ID and an API key each time a protected resource is read or accessed. For more information, see Creating IAM credentials.
• ☐ **Create API keys**	An API key is a unique code passed into an API to identify the calling application or user. You can use platform IBM Cloud API keys that are associated with user identities, and you can create other API keys for service IDs. For more information, see Understanding API keys. If you have configured a Secrets Manager instance as described in Encrypt and protect your data, you can use that instance to dynamically generate a service ID and an API key each time a protected resource is read or accessed. For more information, see Creating IAM credentials.
• ☐ **Learn about IAM policies**	A policy consists of a subject, target, and a role. A policy grants a subject one or multiple roles to a set of resources so that specific actions can be taken. The role defines the level of access that is granted. For more information, see What are IAM policies and who can assign them?

Getting started tasks for streamlining access management with identities, groups, and policies

Get support and other resources

If you experience problems with IBM Cloud, you have several options to get help determining the cause of the problem and finding a solution. If you're logged into your account, you can go directly to the Support center to review product topics and featured FAQs, open or manage a support case, or search community content. Use this checklist to identify support options that are available, including training and tutorials.

Task	Description
• ☐ **What's new in IBM Cloud**	Stay up to date with the new features that are available on the IBM Cloud platform so that you get the most out of your IBM Cloud experience. Check out the what's new documentation for the platform and the announcements blog.
• ☐ **Viewing cloud status**	The IBM Cloud Status page is the central place to find details about all incidents, planned maintenance, announcements, release notes, and security bulletins about key events that affect the IBM Cloud platform and services. Learn more about viewing cloud status.
• ☐ **Review available support plans**	You can choose a Basic, Advanced, or Premium support plan to customize your IBM Cloud support experience for your business needs. The level of support that you select determines the severity that you can assign to support cases and your level of access to the tools available in the Support Center. Learn more about support plans.

• ☐ **Create support cases**	If you experience problems with IBM Cloud, you can use the Support Center to create a support case. You can also create support cases for issues that are associated with access (IAM), billing and usage, account issues, and invoice or sales inquiries. Learn more about creating support cases. In the event you're unable to access your account, please create a case here.
• ☐ **Check out the solution tutorials**	Solution tutorials provide step-by-step instructions on how to use IBM Cloud Log Analysis to manage operating system logs, application logs, and platform logs in the IBM Cloud to implement common patterns based on best practices and proven technologies. See Getting started with solution tutorials.
• ☐ **Technical training and certification resources**	Grow your skills with learning paths. Choose a topic and get started with step-by-step technical training. Visit the IBM Cloud Center for Training.
• ☐ **Give us feedback**	We want to hear from you! You can submit feedback for the IBM Cloud team on the documentation or the console. You can choose from a few different methods to provide feedback. Learn about the ways that you can submit feedback.
• ☐ **Connect and get developer resources**	Learn in-demand skills, build solutions with real sample code, and connect with a global community of developers. Bookmark IBM Developer.
• ☐ **Ask questions and find answers on Stack Overflow**	Ask questions or view responses on Stack Overflow by using the `ibm-cloud` tag or expand to include more related tags.
• ☐ **Find an answer to your FAQ about support**	Visit the FAQ library in the console to get answers about frequently asked questions for getting support.
• ☐ **Working with a Customer Success Manager (CSM)**	A CSM is reserved for those clients who have made a significant commitment to employ IBM Cloud for their cloud-based solution. Assignment takes place by a nomination process with supporting criteria that does vary per geo location. Work through your local IBM Sales team or Inbound Sales for additional details on receiving a dedicated CSM resource.

• ☐ **Submit product ideas and feature requests**	Submit your ideas and feature requests for review and possible implementation by the Product Management team. The Product Management team is responsible for the full lifecycle of an IBM Cloud product offering. Go to the Public Ideas Portal to submit a request.

Getting started tasks for get support and other resources

Running secure enterprise workloads on IBM Cloud

With continuous compliance at the core of IBM Cloud®'s platform, your team has all of the tools at your disposal to securely develop, deploy, and manage your regulated, mission-critical enterprise workloads in the cloud.

For highly regulated industries, such as financial services, achieving continuous compliance within a cloud environment is an important first step toward protecting customer and application data. Historically, that process has been difficult and manual, which placed your organization at risk. But, with IBM Cloud you can work with predefined deployable architectures, automate infrastructure-as-code deployments with *projects*, and integrate automatic security checks into every day workflows to minimize risk.

With IBM Cloud, your whole team from solution architects, compliance managers, infrastructure DevOps teams, to application development teams can use a shift-left approach to identify security risks and exposures early when developing and deploying cloud solutions. This keeps security and compliance at the center of your workflow and organizational culture so that your enterprise can operate in the cloud with confidence. Check out the following video to learn more about how IBM Cloud can help you accelerate and maintain your regulated standards at scale.

View video: Accelerating and Maintaining Your Regulated Standards at Scale

Video transcript

Hi, I'm Mark Meredith a product manager for IBM Cloud.

As a large enterprise your requirements for scale, security, infrastructure, and governance all influence your move to the cloud.

With the tools available through IBM Cloud, you can start and stay compliant, use automation, and ensure deployments are done by using a secure software supply chain, all while managing your resources at scale.

Creating secure, compliant, robust, and scalable application infrastructure can be difficult to set up and costly to maintain.

But, if you standardize application infrastructure where possible and then host many applications on the same infrastructure you can reduce operations costs.

You can do this with deployable architectures and projects on IBM Cloud.

With our growing catalog of deployable architectures you can quickly get started by adding it to a project and customizing it to your needs.

Projects aren't just for configuring your solution though. Projects also enable you to organize resources by what you care about in your organization, whether it's by application or business unit, to get helpful views into cost, availability, and compliance.

Projects have built in governance ensuring that only validated, compliant, and approved code is deployed.

Achieving continuous security and compliance for your regulated workloads is key to your business and ours. With IBM Cloud Security and Compliance Center, you can integrate daily, automatic compliance checks into your development lifecycle to help minimize risk and be audit-ready.

From automating the creation of accounts to authorizing only IBM Cloud projects to deploy resources from approved, compliant deployable architectures, your enterprise can shift left and reduce the possibility of human error that can lead to major security or financial risks.

Visit the IBM Cloud catalog to check out the secure by default deployable architectures that you can add to a project to get started today.

Defining your compliance strategy

By planning and defining your enterprise's goals for running secure workloads on IBM Cloud early in the process, you can reduce the time to production for building infrastructure and applications in the cloud. With IBM Cloud, you can save your business time and money by taking advantage of our automation and standardized best practices as you work. Get started by reviewing our predefined, compliant architectures and control libraries to see how your industry fits in the cloud.

Understanding security and compliance requirements

While you're working through the planning phase of running your enterprise solutions on IBM Cloud, it's important to consider the internal and external standards that your organization must meet. For example, if you are a medical establishment, you most likely need to prove that you are upholding HIPAA requirements. Or, if you are a financial institution you might need to prove that you are meeting NIST standards. With IBM Cloud® Security and Compliance

Center, you can take advantage of predefined control libaries that are designed to help you validate and prove compliance for a group of controls.

Identifying the right infrastructure architecture

Instead of figuring out how to assemble a compliant infrastructure architecture on your own, you can review the deployable architectures that are available in the catalog. IBM Cloud provides automation for the deployment of common architectural patterns that combine one or more cloud resources, known as deployable architectures. Each *deployable architecture* is built and maintained by IBM Cloud experts following IBM Cloud best practices to take the guesswork out of the architecture design process and reduce the time that it takes to deploy to just minutes. Compliance managers and solution architects can review the components of the architecture and the level of compliance that each deployable architecture meets by reviewing the details directly from the catalog detail pages.

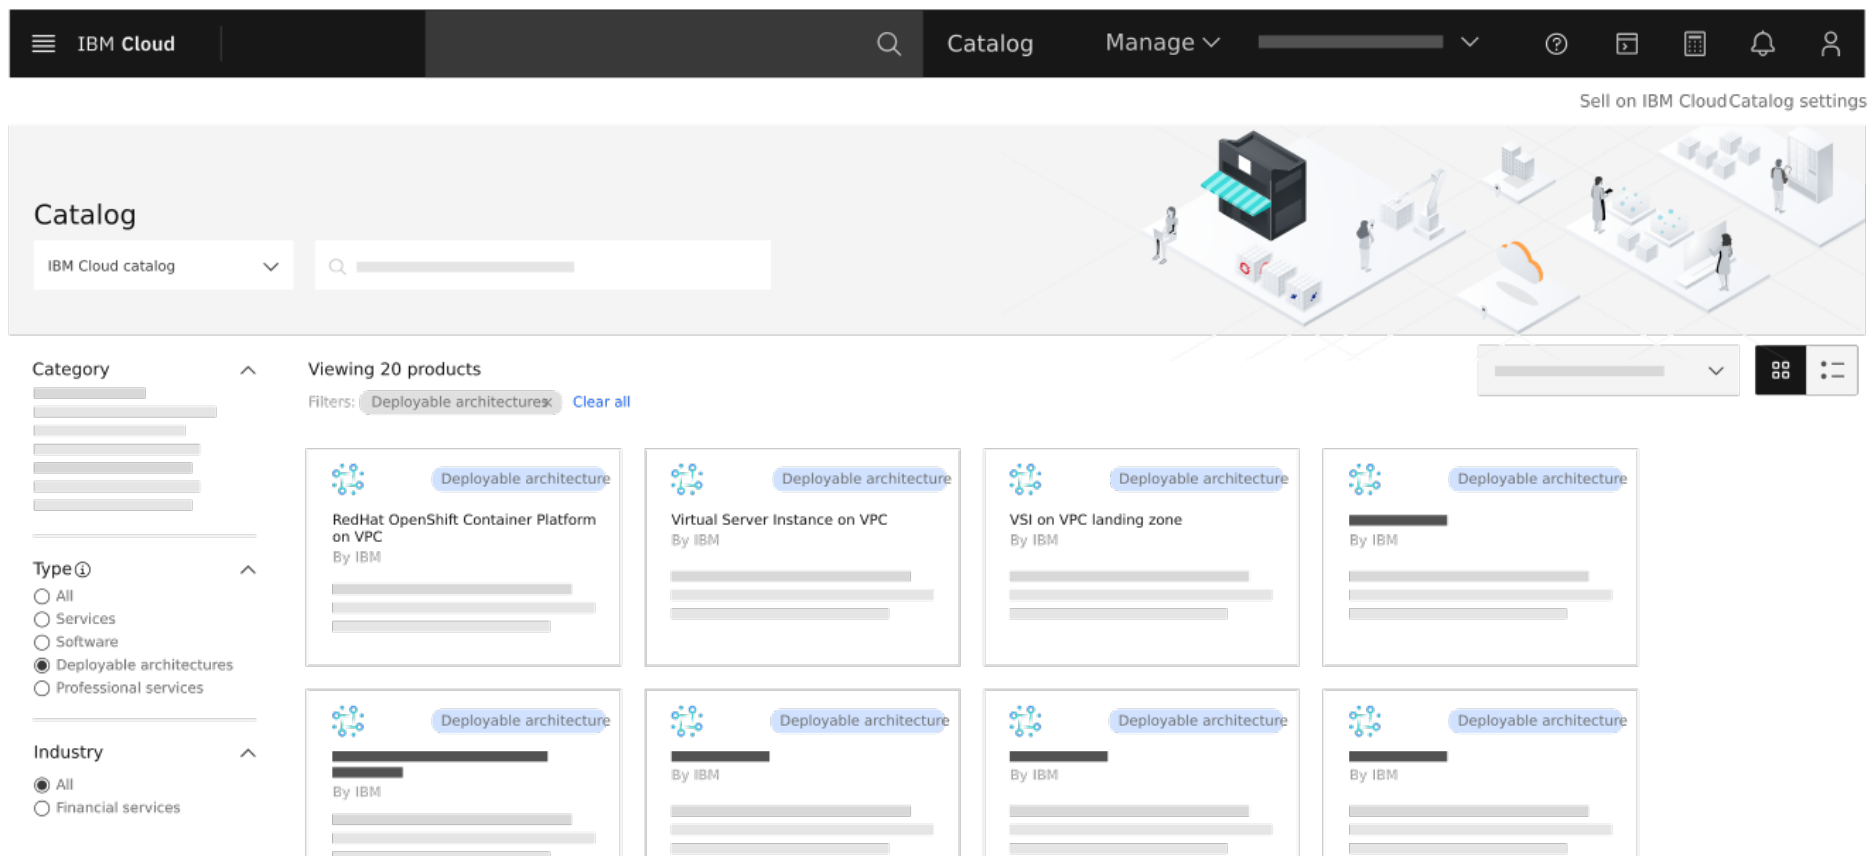

IBM Cloud catalog showing deployable architecture tiles

If you work in a financial institution, you might want to work with IBM Cloud Framework for Financial Services validated products. The framework is designed to build trust and enable a transparent public cloud ecosystem with a focus on the specific features for security, compliance, and resiliency that finanicial institutions require. For more information about the IBM Cloud Framework for Financial Services, see Getting started with IBM Cloud Framework for Financial Services. To get started, you can filter the catalog to view only those products that have been Financial Services Validated. And, even if you're not in the financial industry, you can be assured that these products meet stringent industry regulations for sensitive data and complex workloads.

Configuring and deploying compliant architectures and applications

When your team has evaluated and chosen a deployable architecture, development engineers can use a project to configure it to fit your enterprise's business needs.

A project is a named collection of configurations that are used to manage related resources and Infrastructure as Code (IaC) deployments across accounts. They enable your teams to focus on a shift-left approach by using DevOps best practices to configure, deploy, and monitor deployments. Each project includes tools that scan for potentially harmful resource changes, compliance, security, and cost, as well as tracking versioning and governance. They're designed with an IaC and compliance-first approach that helps to ensure that a project is managed, secure, and always compliant.

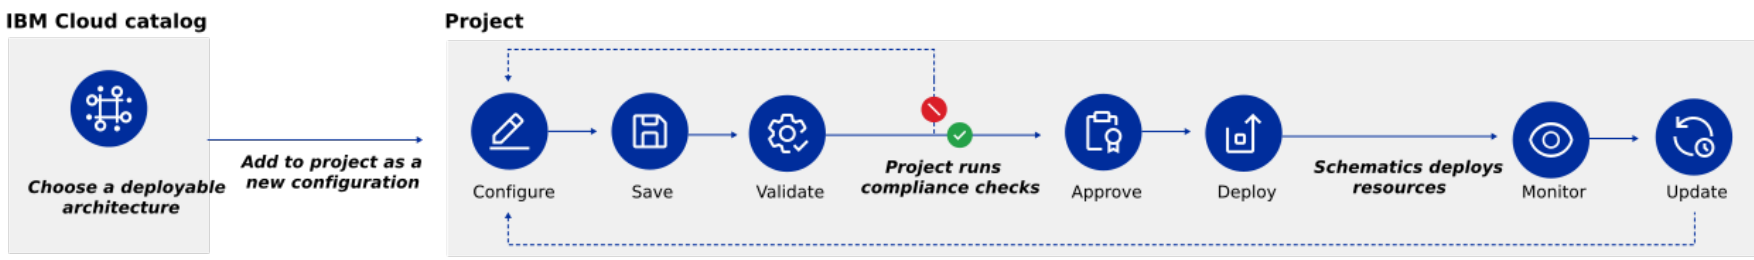

Understanding the projects workflow.

You can customize and extend an IBM Cloud deployable architecture to ensure that the architecture meets your needs. Then, you can leverage private catalogs in your account to make your custom deployable architecture available for only members of your enterprise. After onboarding a customized deployable architecture to your private catalog, users that have access to the private catalog can then configure and deploy it from a project.

Securing your software supply chain

Your team can take advantage of cost-effective software delivery and proactive security practices with IBM Cloud DevSecOps. The DevSecOps Application Lifecycle Management deployable architecture provides a streamlined way to set up continuous integration (CI), continuous development (CD), and continuous compliance (CC) toolchains for secure and agile application development.

By using the DevSecOps Application Lifecycle Management deployable architecture to set up the framework to manage the lifecycle of your application code, you can put security and compliance at the forefront of your development lifecycle and set your team up to implement an automated, shift-left approach to development.

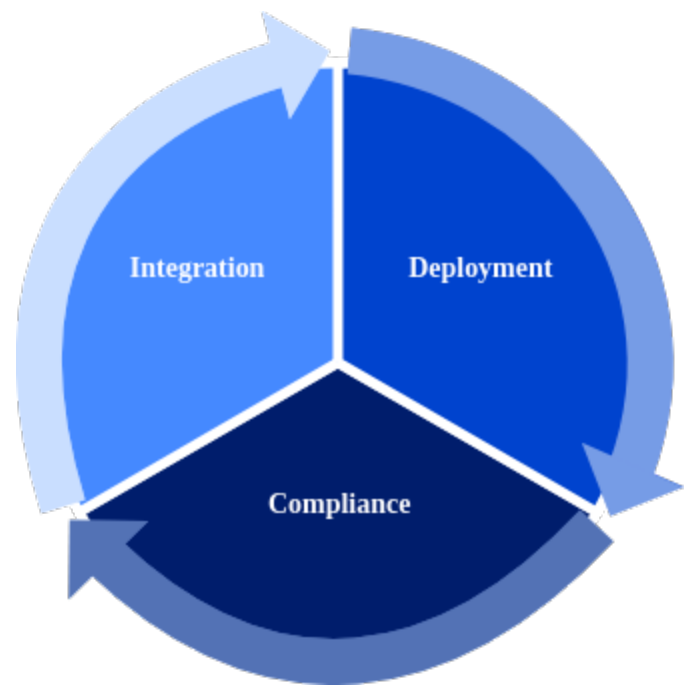

Continuous integration, deployment, and compliance

The CI toolchain provides pipelines that scan your code and build artifacts, catch security issues early, and prevent those issues from being introduced into the product configuration, while collecting evidence that can be used for audits. The CD toolchain checks for required evidence, generates a change request document with an evidence summary and description of changes, and gates deployment based on compliance with your organization's security policies. And, because new vulnerabilities that were not previously known might always be discovered, you can leverage the CC toolchain to continuously assess app security and compliance posture for code that is already in production.

To learn more, see the DevSecOps documentation.

Assessing and maintaining compliance of enterprise workloads

Although IBM Cloud reduces the time and complexity of setting up a compliant enterprise application, you still need to ensure that you're maintaining compliance. To do so, you can use the Security and Compliance Center to run automatic evaluations on your resource configurations. The evaluation results are provided in the dashboard of the Security and Compliance Center or you can get notified of changes. You are able to quickly assess the risk to your organization, fix issues, and generate reports so that you're always audit-ready.

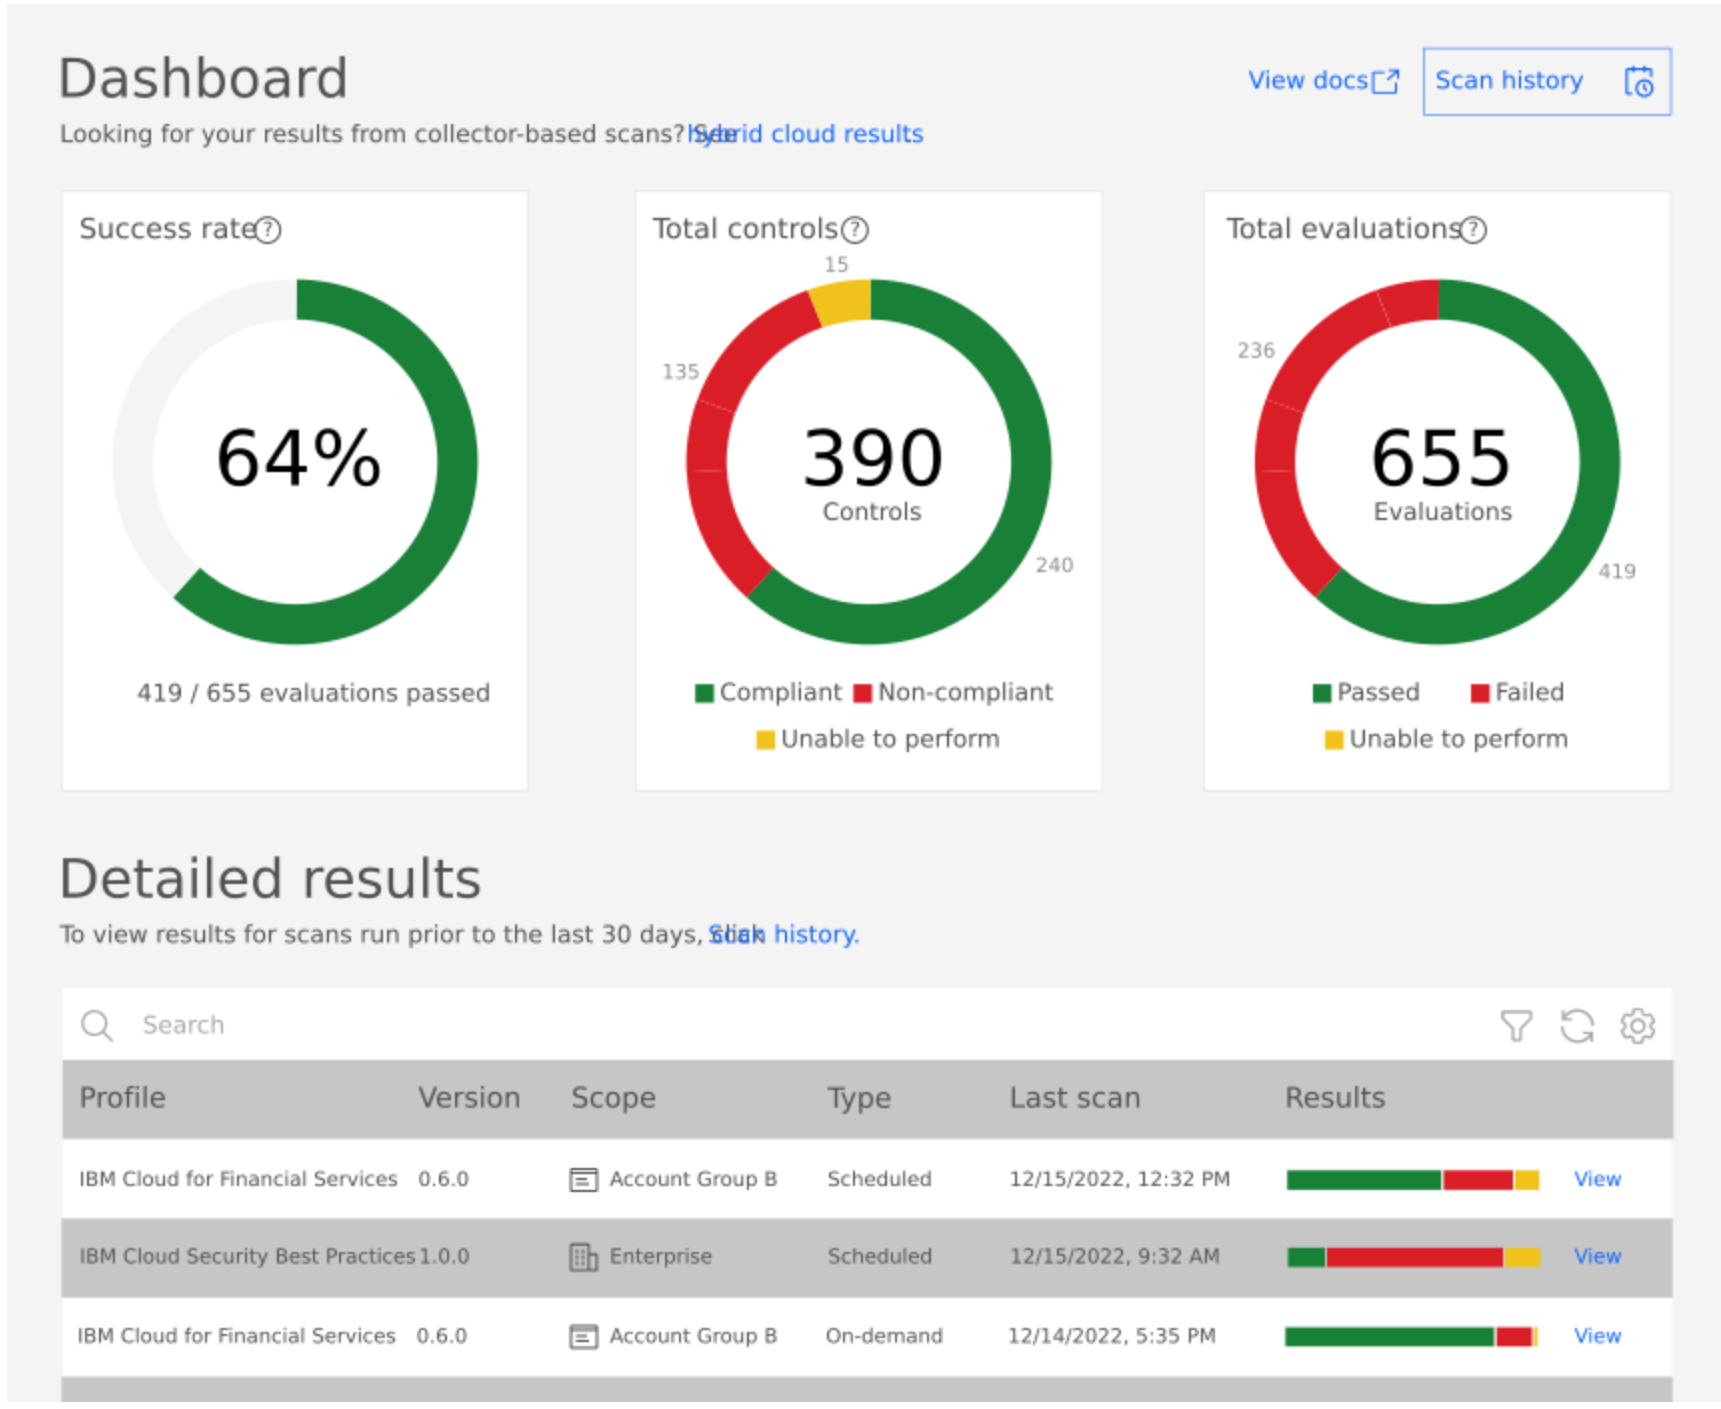

Example Security and Compliance Center dashboard

Also, using DevSecOps CI/CD/CC toolchains can help to automate the evaluation of controls as part of the development process and can block noncompliant changes from being promoted. Managing your application code this way ensures that you have the evidence and change history that is needed to meet the required compliance standards for your industry. For more information about a shift-left approach to deploying your code, see DevSecOps architecture.

Get started with your enterprise workloads on IBM Cloud

With the tools available through IBM Cloud, you can start and stay compliant with automation and ensure that deployments are done by using a secure software supply chain, all while managing your resources at scale. Go to the IBM Cloud catalog to check out the deployable architectures and visit the Security and Compliance Center today to start defining your goals for security and compliance.

For more information about setting up your enterprise, configuring automated deployments by using projects, customizing deployable architectures, and more, see the Enterprise account architecture white paper and Running secure enterprise workloads documentation.

If you're already running workloads on the cloud, but you're not yet taking advantage of an enterprise account structure or tools like IBM Cloud projects to manage related resources and IaC deployments across accounts, you can review the white paper on moving to the enterprise architecture and learn how you can move existing deploy resources to a project.

Understanding compliance in IBM Cloud

IBM Cloud® provides a secure cloud platform that you can trust. IBM Cloud compliance results from a platform and services that are built on best-in-industry security standards, including GDPR, HIPAA, ISO 9001, ISO 27001, ISO 27017, ISO 27018, PCI, SOC2, and others. For more information, see Compliance on the IBM Cloud

Compliance reports

IBM Cloud provides compliance reports for some compliance regulations, for example SOC or PCI, to advise clients. The reports include details of the audit experience and can help users assess and address the high, medium, and low risks identified.

For a description on each security standard and for information on whether to download or request reports, see Compliance on the IBM Cloud.

If you have an IBM Cloud account and want an infrastructure compliance report, go to https://cloud.ibm.com/classic/security/compliancereport/request. Complete the form to receive an email with the requested reports attached.

If you don't have an IBM Cloud account or want a PaaS compliance report, go to https://www.ibm.com/account/reg/us-en/signup?formid=MAIL-wcp. Complete the form and you will be contacted by an IBM representative.

For questions about the compliance program or about the availability of a compliance report, go to Compliance on the IBM Cloud: Next steps for contact options.

General Data Protection Regulation (GDPR)

The GDPR seeks to create a harmonized data protection law framework across the EU and aims to give citizens back the control of their personal data. The GDPR imposes strict rules on those hosting and processing personal data, anywhere in the world.

IBM is committed to providing our clients and IBM Business Partners with innovative data privacy, security, and governance solutions to assist them in their journey to GDPR readiness. Data and data protection are becoming increasingly important to individuals and society. Enterprises must earn the client's trust in their ability to steward information.

IBM Cloud is agile and scalable with built-in data security, and privacy services and solutions that can be consumed on premises or through public cloud. Our comprehensive data security platform helps safeguard sensitive data wherever it resides and provides a full range of data protection capabilities.

Environmental information

IBM Cloud, as a user and as a provider, is environmentally conscious and strives to provide power efficiency and recycling in our data centers. As such, the servers that are put in service within the IBM Cloud comply with Commission Regulation (EU) 2019/424 of 15 March 2019 laying down ecodesign requirements for servers and data storage products (EU Lot 9).

For details, see the following data sheets on our physical hardware in the cloud:

- Rack Mount Server 618U-TR4T+
- Rack Mount Server 6019U-TN4R4T
- Rack Mount Server 5019C-WR-04
- Rack Mount Server 5019S-W4TR
- Rack Mount Server 628U-TR4T+
- Rack Mount Server 6029U-E1CR4T
- Rack Mount Server 2049U-TR4-06-ST031
- Rack Mount Server 2049U-TR4-04-ST031
- Rack Mount Server 2049U-TR4-02-ST031
- Rack Mount Server SR630_1100W
- Rack Mount Server SR630_750W
- Rack Mount Server SR650
- Rack Mount Server 2U_H11DSU
- Rack Mount Server SR650v3
- Rack Mount Server 221H-TNR
- Rack Mount Server SR630v3

IBM Cloud Security and Compliance Center

For highly regulated industries, such as financial services, achieving continuous security within a cloud environment is an important first step toward protecting customer and application data. Historically, that process has been difficult and manual, which placed your organization at risk. But, with the IBM Cloud® Security and Compliance Center you can integrate automatic security checks into every day workflows that are designed to minimize risk. By monitoring for risks, you can identify security vulnerabilities and quickly work to mitigate the impact and fix the issue.

You can acheive continuous security and compliance within the IBM Cloud platform with the Security and Compliance Center. You can view your security and compliance postures from a unified dashboard. Learn more about Getting started with Security and Compliance Center.

IBM Cloud Framework for Financial Services

IBM Cloud Framework for Financial Services is designed to build trust and enable a transparent public cloud ecosystem with the specific features for security, compliance, and resiliency that financial institutions require. The IBM Cloud Framework for Financial Services is a standard set of controls, controls guidance, and architectures informed by global regulatory requirements for cybersecurity, data security, and risk management, with ongoing governance by an industry council and promontory to ensure currency with new and changed regulations.

IBM Cloud Framework for Financial Services currently applies controls (US NIST 800-53 with IBM financial services guidance) to IBM Cloud services, IBM software, and third-party ISV and SaaS providers that provide a common control approach that can be mapped to regulatory guidelines worldwide. This solution platform and ecosystem program is built on an industry-informed framework of controls, architectures, and operations that mitigates systemic risk in using the IBM public cloud for mission-critical workloads with client-sensitive data. For more informations, see Getting started with IBM Cloud for Financial Services and how to enable your account to use Financial Services Validated products.

EU support for your account

You can add an extra layer of protection to your data in the European Union. Turning on EU support for your account ensures that any support ticket that you create is sent to IBM support teams in Europe. An EU supported account is useful if, for example, you use resources to process personal data for European citizens. For more information, see Enabling EU support for your account.

HIPAA support for your account

The US Health Insurance Portability and Accountability Act (HIPAA) and the Health Information Technology for Economic and Clinical Health (HITECH) Act define standards for handling electronic healthcare transactions and information. If you or your company is a covered entity as defined by HIPAA, you must enable the HIPAA Supported setting if you run sensitive workloads that are regulated under HIPAA and the HITECH Act. By using this setting, you can filter on HIPAA Enabled services in the catalog, indicate to IBM that your account stores protected health information (PHI), and digitally accept the IBM Business Associate Addendum for covered entities. For more information, see Enabling HIPAA support for your account.

Understanding data residency in IBM Cloud

IBM Cloud's global network of locations provides you with the flexibility of choosing where you want to run your workloads.

For regional and zonal services, you request instances of the service to be deployed in a specific region in accordance with your specific geographic requirements. IBM Cloud ensures that content that is provided by you and your workload (as defined in the IBM Cloud Service Agreement) is stored and processes locally in the selected region location. For a complete list of the locations where IBM Cloud services are available see Service and infrastructure availability by location.

IBM Cloud also maintains internal encrypted backups of the customer content within the same geography where the regional or zonal service is located for recovery in case of data corruption or a major data center disaster.

About your business contact and account usage information (as defined in the IBM Cloud Service Agreement), also called client's metadata, IBM Cloud stores and processes them where the control planes of the regional and global services are located.

- Regional services typically have control planes that are located in the same region where you selected for the service except for the services indicated in Services with global control planes.
- Global services control planes locations are indicated in Global platform services.

For a complete list of data attributes that are stored and processed by each single IBM Cloud service you can refer to the documentation of the API of each service in the API and SDK reference library.

All data in transit is encrypted. Only TLS 1.2 and 1.3 are supported in IBM Cloud with TLS 1.1 and below explicitly disabled to prevent downgrading to a vulnerable version of the protocol.

IBM Cloud data privacy processing processes and procedures are documented within the IBM Cloud DPA. This Data Processing Addendum (DPA) and its applicable DPA Exhibits apply to the Processing of Personal Data by IBM Cloud on behalf of Client (Client Personal Data). The processing of Personal Data is subject to the General Data Protection Regulation 2016/679 (GDPR). It is also subject to any other data protection laws that are identified at Data Protection Laws in order to provide services (Services) according to the Agreement between Client and IBM Cloud. The IBM Cloud DPA can be found at Data Processing Addendum.

In addition to the DPA, the cloud services provide DPA exhibits that can be found on the IBM Cloud Terms site.

How do I know that my data is safe?

Designed with secure engineering practices, the IBM Cloud® platform provides layered security controls across network and infrastructure. IBM Cloud focuses on protection across the entirety of the compute lifecycle, which includes everything from the build process and key management to the security of data services. IBM Cloud also provides a group of security services that can be used by application developers to secure their mobile and web apps. These elements combine to make IBM Cloud a platform with clear choices for secure application development.

In addition to our own diligence in creating and operating a secure cloud, IBM® also engages many different firms to assess the security and compliance of our cloud platform. For more information, see IBM Cloud compliance programs for a detailed list of certifications and attestations.

IBM Cloud ensures security readiness by adhering to security policies that are driven by best practices in IBM for systems, networking, and secure engineering. These policies include practices such as source code scanning, dynamic scanning, threat modeling, and penetration testing. IBM Cloud follows the IBM Product Security Incident Response Team (PSIRT) process for security incident management. See the IBM Security Vulnerability Management (PSIRT) site for details.

In addition to the regular penetration testing conducted by IBM and our partners, you can conduct penetration testing of your VPC or Classic Infrastructure resources on IBM Cloud. Prior authorization to do so is not required by IBM Cloud. IBM Cloud customers under an active NDA can request a copy of a penetration testing executive summary by opening a support case.

For more details about security for your applications and environments in IBM Cloud, see Security for IBM Cloud.

Quantum safe data protection

Quantum computing promises to solve complex problems that even the most powerful computers can't solve today. At the same time, there is risk that data protected by public key cryptosystems could be recorded today and decrypted years later by using a Cryptographically Relevant Quantum Computer (CRQC). The IBM Cloud platform provides a secure, reliable, and cost-effective cloud computing environment that's tailored to your specific business needs. With security at the core, IBM Cloud offers capabilities that you can use to integrate your apps and tools with the required level of data protection.

With a focus on the prioritization of workload protection, IBM Cloud includes the following core features that you can leverage at rest and in transit for inbound and outbound traffic.

Data at rest

Key management, using Key Protect and Hyper Protect Crypto Services, supports large key sizes that are considered quantum safe for data encryption (Data Encryption Key) and envelope encryption (Key Encryption Key). For more details, see Protecting data with envelope encryption and Bringing your encryption keys to the cloud.

Data in transit

Key Protect supports quantum safe enabled TLS connections through a hybrid method that combines Quantum Safe Cryptography and current ECC algorithms. Key Protect uses the Kyber algorithm with NIST evaluation round three parameters. See Introduction to Quantum-safe Cryptography in TLS for more details.

Secure your outbound data with post-quantum support on IBM Cloud Internet Services. See Bringing post-quantum cryptography to IBM's edge for more details.

For cloud native apps, TLS connections are quantum safe enabled with a custom ingress controller for IBM Cloud Kubernetes Service and a custom router for Red Hat OpenShift on IBM Cloud. See Protecting apps on IBM Cloud with Quantum Safe Cryptography for more details.

Authentication

A high-security version of round 2 Dilithium digital signatures in Hyper Protect Crypto Services is primarily used for data integrity, authenticity, and non-repudiation. See Post-quantum cryptography support for more details.

What's a Customer Success Manager?

An IBM Customer Success Manager (CSM) is a technical architect who works as part of your team to accelerate your time to value and help you achieve your business outcomes as you use IBM Software and IBM Cloud services.

CSMs work with IBM customers to:

- Identify and implement use cases
- Review and co-create MVPs
- Design and validate solution architectures, enterprise account and user governance models
- Guide MVPs into production and scale with Business Partners, IBM Expert Labs, and IBM Consulting teams
- Optimize resource usage and cost management
- Review security and compliance requirements and implementations
- Ensure effective and efficient operational models

Additionally, CSMs [1] engage product teams for roadmap previews, engage renewal teams for a seamless renewal process, provide demos for business and technical sponsors, help transfer knowledge of best practices, create integrated learning plans, and run business reviews to ensure alignment on outcomes.

At the end of the day, the CSM's success is your success.

Explore some of the products that CSMs support

Product	Description
Asset Management	IBM Maximo Manage is a fully integrated asset management platform that uses advanced analytic tools and IoT data to improve operational availability, extend asset lifecycles and optimize performance.
Supply Chain	Supply chain resiliency starts with a business network that handles peak events with ease and detects potential issues before they occur.
Data Fabric	Data fabric is an architecture that facilitates the end-to-end integration of various data pipelines and cloud environments through the use of intelligent and automated systems.
Business Automation	IBM Cloud Pak for Business Automation is a modular set of integrated software components, built for any hybrid cloud, designed to automate work and accelerate business growth.
Security	IBM Cloud Pak for Security can help you gain deeper insights, mitigate risks and accelerate response.
IBM Cloud	IBM Cloud provides solutions that enable higher levels of compliance, security, and management, with proven architecture patterns and methods for rapid delivery for running mission-critical workloads.

Learn how our customers are finding success

The following use cases are examples of CSMs collaborating with IBM customers.

NCHC uses AIOps to improve public network services and proactively prevent outages.

> In Taiwan, where the pandemic response has been exceptionally effective at limiting outbreaks and death, the National Center for High-performance Computing (NCHC) helps accelerate research and innovation nationwide by providing access to supercomputers and analytics and by facilitating nationwide networks for data sharing and collaboration. Read the case study.

Travelping drives connected cars into the data fast lane with cloud-native technologies.

> EU data protection regulations and network latency issues can stall advancements in connected vehicle performance. To help auto manufacturers gain real-time Internet of Things (IoT) insights, Travelping used IBM Cloud® infrastructure products and its own Kubernetes deployment to create a cloud-native solution that transports data to the vehicles. Read the case study.

Karnataka Bank creates single integration platform for its data-driven transformation.

> Financial institutions are constantly operating in a complex, competitive and compliance-driven financial environment where making timely and data-driven decisions is key for efficient growth and effective risk management. By leveraging tools such as data lake, advanced analytics, and

accessible reporting, a bank can establish baseline performance metrics across applications, deposits, and loans. Read the article.

Pandora offers customers a personlized digital expereinces.

With IBM Sterling Order Management on Cloud, Pandora is digitally transforming online sales, order management, fulfillment and returns, to offer customers personalized digital interactions similar to the in-store experience. Sterling platform in Europe was online in less than five months, with excellent results. Read the case study.

---

1. CSMs are assigned to customers at IBM's discretion. ↩

Release notes for the IBM Cloud platform

Stay up to date with what's new in IBM Cloud® and the latest feature releases that are available on the platform so that you get the most out of your IBM Cloud experience.

December 2024

17 December 2024

Experimental: Customize a deployable architecture from the catalog

> You can now customize a deployable architecture by adding optional components in the IBM Cloud console. Add the Cloud automation for Secrets Manager deployable architecture to a project to get started.
>
> For more information about onboarding a deployable architecture and specifying optional components that work with it, go to Optional and swappable components for deployable architectures.

13 December 2024

Service ID limits are increased with groups

> With Service ID groups, your account can now contain up to 100,000 service IDs, with up to 2,000 service IDs in each group. Existing service IDs were added to the default group. For more information, go to Creating and working with service IDs .

November 2024

21 November 2024

Experimental: Connect your project to a Git repository

> You can now connect a project to a Git repository to save configurations there. By doing so, you can use your repository and the CI / CD tools of your choosing to automate pipelines on configurations. Connecting a project to a Git repository works best with a new project that doesn't contain any configurations. This is an experimental feature that is available for evaluation and testing purposes and might change without notice. For more information, go to Integrating a project with a Git repository .

12 November 2024

The AI assistant is generally available

> The AI assistant is available to help answer your questions about working in IBM Cloud. In addition to being able to launch the AI assistant in the IBM Cloud console, you can also ask questions from the IBM Cloud CLI. For more information, see Getting help from the AI assistant.

October 2024

17 October 2024

Consolidated and simplified IBM Cloud platform

> IBM Cloud is bringing simplification to your door. The IBM Cloud platform recently consolidated and simplified services and features. Services and areas of the console are now unified into the following hubs: Infrastructure, Containers, Automation, Databases, Observability, and Security with the goal to bundle together related services to make it easier to find, deploy, and use them. Users can now easily search for new services to accelerate

business-critical workloads and better manage what is already running to help ensure their business stays on top of its velocity. For more information, see Navigating the IBM Cloud console.

Improving the account and enterprise documentation experience

Easily navigate enterprise information with the latest update to IBM Cloud docs. Now, you can browse key information about enterprises in one location. The documentation that supports enterprise accounts is moved to Getting help from the AI assistant.

Changes to IBM Cloud projects API method to list all deployed resources

A change was made to the `project.config.retrieve-resources` API method to improve its performance. This change can cause some of the resources to temporarily disappear from the **Resources** tab in your configurations. The resources were not destroyed. Redeploy the configurations to see the resources again.

Experimental: Project Resources tab

Manage the resources in your project in the new **Resources** tab. You can add existing resources to your project without using a configuration. To get started, see Adding existing resources to a project.

May 2024

31 May 2024

Onboarding software with usage-based pricing plans

In addition to free and bring your own license plans, you can now add usage-based pricing plans to software onboarded with the virtual server image delivery method, offering it as a paid integrated product in the IBM Cloud catalog. For more information, see Adding a usage-based plan and Onboarding a virtual server image for VPC with a plan.

After adding a plan to your software in Partner Center, you can also manage it in your private catalog. This includes adding features, changing the plan's state, deprecating it, or updating plan details. For more information, see Managing software plans in catalogs.

18 May 2024

Google login and registration

Google is now available for IBM Cloud login and IBM Cloud registration. As a non-federated user, you can use your existing Google credentials to securely log in to IBM Cloud or create a new IBM Cloud account. Learn more.

06 May 2024

Experimental: Stack deployable architectures by using the console

Create more complex end-to-end solutions, without editing Terraform, by stacking deployable architectures. To get started, see Stacking deployable architectures.

April 2024

10 April 2024

Discover your software resources by category

Software resources that are deployed in your account are now able to be discovered in their associated catalog category. To quickly find your product, you can still take advantage of the console's filtering capabilities.

03 April 2024

Projects API v1.0.0

Version 1.0.0 of the projects API is now available. To learn more about the changes that are included in this release, see the Projects API change log.

Experimental: Stacking deployable architectures

You can now stack deployable architectures together to create more complex end-to-end scenarios by using either the API or the CLI. To learn more about stacking deployable architectures, see What is a deployable architecture?

March 2024

21 March 2024

Increased limits for enterprises

Easily scale your organization and workloads with increased limits on the number of accounts and account groups that you can have in your enterprise. You can have a maximum of 1000 accounts and 500 account groups. For more information, see Enterprise limitations.

05 March 2024

Upgrade your Pay-As-You-Go account

You can now create an enterprise from Pay-As-You-Go accounts that signed up with a credit card on cloud.ibm.com. For more information, see Billing FAQs.

February 2024

20 February 2024

Organize existing resources in your project

Now, you can use a project to organize resources across accounts. By doing so, your resources remain deployed in their respective accounts, but you gain an at-a-glance view of those resources in your project. Organizing resources is experimental and might change without notice. For more information, see Organizing existing resources by using a project.

06 February 2024

Automatically detect drift in your project

You can now run a daily scan to detect any changes between your configured architecture and the actual state of your deployed resources. To fix drift, you can override the changes or adopt them. Automatic drift detection is an experimental feature and might change without notice. For more information, see Managing Drift.

Add references in your project

It's now easier to link architectures to one another by adding references to inputs or outputs. Add the reference as an input in an architecture that you're configuring, and the value is automatically used for deployment. For more information, see referencing values.

Duplicate configurations in your project

Now, you can duplicate an architecture configuration from the Configurations tab in a project. By doing so, you can quickly create a second version of your configuration with all of the input values that were provided in the original configuration. From the **Configurations** tab, select the **Actions** icon

for the configuration you want to duplicate and click **Duplicate**.

December 2023

07 December 2023

Limit access with resource attribute-based conditions

IBM Cloud IAM is excited to give customers the ability to grant access based on multiple resource attribute-based conditions. With this ability, you can create a single policy with various conditions instead of individual access policies. For more information, see Limiting access with resource attribute-based conditions.

November 2023

17 November 2023

Protect your private catalog and private products

Add an extra layer of protection to your private catalogs with context-based restrictions. Users can view the catalogs that you protect only if they satisfy your rule.

14 November 2023

Move deployed software into a deployable architecture and manage it by using a project

Use the IBM Cloud Catalogs management CLI plugin and the Project CLI plugin to convert deployed software or an existing Schematics workspace and its resources into a format that is compatible with using projects to manage future deployments. For more information, see Moving resources from a Schematics workspace into a project and the white paper on Moving to the enterprise architecture.

09 November 2023

Onboarding virtual server images with Terraform is deprecated

As of 9 November 2023, the option to onboard virtual server images with Terraform is deprecated. If you onboarded software by using this delivery method before this date, your customers can still install the VSIs, but you can't provide any version updates. After 29 March 2024, onboarding virtual server images with Terraform is no longer supported as a delivery method, which means that no new virtual server images with Terraform can be onboarded. Existing VSIs in the IBM Cloud catalog will be available to use, but to take advantage of version updates and ensure continued support, onboard virtual server images for Virtual Private Cloud directly. For more information, see Onboarding a virtual server image for VPC.

06 November 2023

Introducing environments and general improvements to projects API and CLI

You can now group configurations together within a project by using environments. Environments can also contain properties like input values that are automatically added to configurations. This release also includes updates to the projects API and CLI. The commands are normalized so they follow a regular pattern, and terminology is now aligned across our UI, CLI, and API.

For more information on environments, see Controlling deployments by using environments. For more information on updates to the projects API, see the projects API change log. For more information on updates to the projects CLI, see the projects CLI change log.

October 2023

30 October 2023

Selecting currency

The option to select local currency and USD is now available during IBM Cloud registration for users in Canada, United Kingdom/Ireland, Italy, Spain, France, Germany, Austria, Switzerland, Belgium, the Netherlands, Luxembourg, Denmark, Norway, Sweden, Finland, Australia, and New Zealand.

26 October 2023

Symantec VIP 2FA is deprecated

If you are currently using Symantec VIP as your external authentication method, you can continue to use it and add subscribers as needed until 26 October 2024. If you are not already using Symantec VIP, then you can't purchase Symantec VIP to use as your external authentication method as of 26 October 2023.

All customers currently using Symantec VIP must migrate to IBM Cloud MFA by 26 October 2024. For more information about how easy is it to migrate to this no-cost MFA option, see Migrating to IBM Cloud MFA.

19 October 2023

Partner Center Build & Grow is deprecated

As of 19 October 2023, Partner Center Build and Grow is no longer supported. The new simplified and unified build experience has moved to the IBM Partner Plus website, where you can start enrolling. For more information, go to the Partner Plus website and enroll by clicking **Join IBM Partner Plus**.

September 2023

20 September 2023

Support for onboarding third-party deployable architectures in Partner Center

You can now onboard deployable architectures in Partner Center and publish them to the IBM Cloud catalog. A deployable architecture is cloud automation for deploying a common architectural pattern that combines one or more cloud resources that are designed for easy deployment, scalability, and modularity that solves a customer-defined problem. For more information, see Getting set up to sell deployable architectures.

6 September 2023

Cross-account product sharing

You can now share products from your private catalog with other accounts without publishing them to the IBM Cloud catalog. For more information, see Sharing private catalog offerings. The administrator of the account that you want to share your product with must accept the share request to opt in to accessing the offerings that you share. For more information, see Accepting share requests for private catalog offerings.

New CLI commands for cross-account product sharing

New `ibmcloud catalog account` commands are available for cross-account product sharing. For more information, see the Catalogs management CLI plug-in documentation.

Expanded integration between catalogs and projects

You can now link a catalog with a project. A successful validation of a linked project will validate the associated version as well. For more information, see the new section Setting up a target account in the docs.

Catalog and onboarding support for IBM Z® when deploying VSI images

Partners and IBM Cloud users can now onboard Virtual Server Images with IBM Z deployment support. Navigate to your private catalog to get started adding this image. For more information, see Onboarding software to your account.

5 September 2023

Centrally administer your multi-account environment

IBM Cloud is excited to announce that you can now centrally administer access and security settings in your organization with Enterprise-managed IAM. By using enterprise IAM templates, you can quickly create and manage access groups and trusted profiles across accounts and easily ensure that all new accounts inherit the same settings. Leveraging Enterprise-managed IAM can help your organization save time and effort, reduce the risk of unauthorized access, stay compliant with industry regulations and better protect your data. For more information, see Best practices for assigning access in an enterprise.

1 September 2023

IBM Cloud January 2024 price changes

Effective 1 January 2024, prices for some services will be increasing. For more information, see January 2024 price changes.

July 2023

25 July 2023

Projects integration with IBM Cloud Schematics and general improvements

Projects are integrated more seamlessly with Schematics, as users can now view the Schematics logs directly from their project during validation and deployment. This release also includes system improvements to deleting a project. A recent update to the Code Risk Analyzer might result in fewer validation errors for deployable architectures that include Event Streams resources.

11 July 2023

New multizone region (MZR) available in Madrid

Resources can be deployed to a new MZR in Madrid, Spain. There are three data centers available in the `eu-es` region: MAD02, MAD04, and MAD05. For more information about location availability, see Locations for resource deployment.

6 July 2023

Projects API response models update

The response models for all methods now enforce a lower snake case format in state values. This format needs to be expected in the response when you are calling the project and configuration endpoints. This update is a breaking change. For more information, see the change log.

June 2023

6 June 2023

New event notifications for projects

You can now send out notifications when a deployment is complete or when resources are successfully destroyed. For more information, see the available events for projects.

Support for access management tags in the projects UI

You can now add, edit, and delete access management tags directly in the projects UI. Previously, you could manage these tags only by using the command-line interface (CLI) or API.

April 2023

25 April 2023

Customized compliance validation

In addition to using default controls, you can specify a Security and Compliance Center attachment to validate deployable architectures that you're configuring in a project. For more information, see Configuring the architecture.

Support for deployable architectures onboarded from private or public GitHub repositories

You can now deploy architectures that are onboarded to the catalog from either a private or public GitHub repository. Previously, only deployable architectures onboarded from a public GitHub repository were supported.

17 April 2023

Check out deployable architectures in the catalog

IBM Cloud provides deployable architectures in the catalog, which are products that provide automation for the deployment of common architectural patterns that combine one or more cloud resources and are designed for scalability and modularity. Go to the catalog, and filter by Deployable architectures to review the growing catalog of options. For more information, see Identifying the right infrastructure architecture.

IBM Cloud projects for automated IaC deployments

You can now configure, deploy, and monitor deployments by using DevOps best practices with projects. By using projects, you can manage Infrastructure as Code (IaC) at scale and across accounts to ensure that the configuration is always valid, secure, and compliant. Learn more about IaC deployments with projects.

Onboard customized deployable architectures for your enterprise users

You can customize IBM Cloud deployable architectures to meet your enterprise's needs, and then leverage private catalogs to make only approved and compliant architectures available for your enterprise developers to deploy. For more information, see Customizing an IBM Cloud deployable architecture.

12 April 2023

Specify language support for community-supported products

If you offer a third-party product that is supported by open source communities, you can select the languages in which support is provided when you're onboarding or managing your product in Partner Center. By providing support in multiple languages, your customers can communicate their issues in their preferred language and get support more efficiently. For more information, see Defining your support experience.

10 April 2023

View email notification history

You can use the Communication history page in the console to check the status of all email notifications that are sent you to verify if the emails are being delivered successfully. You can also view the last 90 days of IBM Cloud email history, which can help save you time troubleshooting any delivery issues without needing to contact IBM support. For more information, see Checking the delivery status of email notifications and viewing email history.

04 April 2023

Customize the look and feel of your private catalog

You can enhance the appearance of your private catalog to match your brand by adding a custom banner image to your private catalogs. For more information about personalizing your private catalog, see Branding your private catalog with a custom banner .

Add a custom provider name for your private products

You can make it easier for users to search for your products added to a private catalog by specifying a custom provider name. Adding a custom provider name for your private products can help users find them quickly by using the Provider filter in the catalog. For more information, see Providing catalog entry details.

March 2023

29 March 2023

Generate a report on the MFA status of account users

Users that don't meet MFA requirements leave your account vulnerable. You can now identify the users in your account that don't satisfy your MFA requirements. For more information, see Identifying a user's MFA status .

An extra layer of security for users that don't use MFA

IBM Cloud recommends enabling multifactor authentication (MFA) for all users in your account, but some automation scenarios might require you to exclude specific users from your MFA requirement. For users that are excluded from MFA, you can make access more secure by disabling CLI logins with only a username and password. This way, you require an API key to log in to the CLI or users can log in with `--sso`. For more information, see Disabling MFA.

February 2023

22 February 2023

User-specific MFA

You can now enforce user-specific multifactor authentication (MFA) requirements that differ from the account default MFA setting. After you update the MFA requirement for an individual user, view a list of users that have unique MFA requirements in the account by going to Manage > Access (IAM) > Settings > Authentication. For more information, see Enabling MFA for an individual user .

20 February 2023

Download and read IBM Cloud docs offline

While digital is still the primary delivery target for IBM Cloud docs, you can now print or download and read a doc set offline. Go to the menu next to the doc set title in the navigation and click **View as PDF**.

January 2023

27 January 2023

New services integrating with context-based restriction

Services continue to integrate with the IBM Cloud platform's context-based restrictions feature, including IBM Cloud® Virtual Private Cloud (VPC) Infrastructure Services. To see a full list of services that can use context-based restrictions to define and enforce access restrictions on their resources, see Services integrated with context-based restrictions.

25 January 2023

Time-based conditions in IAM access policies

IBM Cloud IAM is excited to give customers the ability to set access controls based on a specified time and date. You can now create policies that grant employees access to a resource during only their working hours, or grant automated processes temporary access for a specified duration. Implementing such limitations helps you to apply the principle of least privilege for assigning access and reduces the opportunity for attack in the event of a security breach. For more information, see Limiting access with time-based conditions.

IAM Policy Management API V2 release

A new version (`v2`) of the IAM Policy Management API is now available. This version adds a new JSON schema to support a conditional policy construct and several time-based comparison operators. These operators provide the capability to restrict access based on time and date. With time-based access control, customers can establish granular policy enforcement based on a specified time period. For more information, see the IAM Policy Management API change log

09 January 2023

Save and share multiple estimates by using the IBM Cloud Cost Estimator

The IBM Cloud Cost Estimator now allows a customer to save multiple estimates to their account and share those estimates with users that belong to the same account. For more information, see Estimating your costs.

December 2022

15 December 2022

Checklist for getting started on IBM Cloud

A new checklist is now available in the IBM Cloud docs for Getting started on IBM Cloud. This checklist outlines the tasks for you to complete to accelerate your journey to cloud by guiding you through your account setup and organization of resources.

October 2022

10 October 2022

A new way to provide the IBM Cloud docs team your feedback

We want to hear from you! Let us know about your experience with the IBM Cloud docs or API docs by submitting your feedback from the **Feedback** button. You can tell us what you love and even what you'd like to see improved, so we can work together to ensure that you always enjoy the IBM Cloud documentation.

07 October 2022

Setting an alternative account owner

As the owner of an account with classic infrastructure, you can set a trusted profile as the alternative account owner. An alternative account owner ensures that you always have a secure way to manage account ownership if your account owner leaves your organization or isn't available.

04 October 2022

Catalog integration with Virtual Private Cloud custom images

ISV Partners and IBM Cloud customers can now import custom VPC images directly to their account or enterprise, or sell on the IBM Cloud catalog without depending on Terraform. For more information, see Onboarding a virtual server image for VPC.

September 2022

29 September 2022

Enabling IBM Cloud Event Notifications for the notification distribution list

You can now easily add Event Notifications instances to the notification distribution list and receive account-wide IBM Cloud notifications. With Event Notifications, you can choose to deliver your notifications to different destinations, including email, SMS, or webhooks. When an event of interest occurs on the IBM Cloud platform and an event is generated, the notification distribution list communicates with a connected Event Notifications instance to forward a notification to the supported destination.

For more information, see Enabling Event Notifications for the notification distribution list.

26 September 2022

Identifying inactive policies

To reduce the number of policies in your account and keep only the minimum access that is necessary for each user, you can now identify the infrequently used access policies on the Inactive policies page in the console. You can determine whether to remove the inactive policies, or in some cases, you might expect an infrequently used policy. For more information, see Managing inactive policies.

Exporting user access reports

Make sure that users have only the access that they need. Export an access policy report for any user in your account to view all of the access policies that they are assigned. For more information, see Exporting user access policy reports.

22 September 2022

A new way to provide IBM Cloud your feedback

We want to hear from you! Let us know about your experience with the IBM Cloud console by submitting your feedback from the **Help** > **Send feedback** option in the console menu bar. You can tell us what you love and even what you'd like to see improved, so we can work together to ensure that you always enjoy developing on IBM Cloud.

21 September 2022

Getting help in the console

You can now find the documentation and Support Center help options from a single menu item in the IBM Cloud console menu bar. Check out the new **Help** option from the console menu bar that now includes your links for Docs, the Support Center, a form to give us feedback, and available guided tours.

16 September 2022

Attaching tags on service IDs

In addition to tagging resources, you can now attach user tags and access management tags on service IDs. User tags help you group service IDs for usage reports, and you can use access management tags to control access to your service IDs. For more information, see Working with tags.

Specifying a user onboarding strategy

If you're Enabling and connecting your identity provider, you can now specify how you want to onboard users to the account upon first-time login. This way, you can add each user to your account when they log in the first time, add users to your account only if they log in and don't select a trusted profile, or never add users to your account and provide access only by using trusted profiles. For more information about trusted profiles, see Creating trusted profiles.

02 September 2022

Increased policy limit

The shared limit for IAM policies and context-based restrictions has increased from 2010 to 4020 so that you don't need to request an increase to the limit in your account as you grow your organizational capacity. For more information, see Known issues and limitations.

August 2022

31 August 2022

Choose from dark or light themes in the console

Enjoy a customized experience in the console that matches your preferences. By default, the console uses the mode that you already have set for your device. We'd like you to be able to experience IBM Cloud using a theme of your choosing. Use the Avatar icon from the console menu bar to find the options for changing your theme.

July 2022

21 July 2022

Monitoring context-based restrictions

To help you predict how context-based restrictions might affect users, applications, and workflows, IBM Cloud is excited to release report-only mode for rules. You can enable context-based restrictions during creation, or choose to set the rule to report-only mode. Using Activity Tracker, you can monitor the impact of enabled rules, or view report-only rules to see how they will affect your users, applications, and workflows without enforcing the rule. For more information, see Monitoring context-based restrictions.

08 July 2022

Streamlined process for updating published support information in Partner Center

You can update support information for your published products by editing your already approved support information and then requesting an

approval. After your changes are approved, you can publish the updates. For more information, see Updating your product's support information.

June 2022

23 June 2022

Onboarding software to sell on IBM Cloud by using the API

To sell your products on IBM Cloud, you can now onboard and publish by using the Partner Center Sell API, in addition to using the Partner Center experience in the console.

22 June 2022

Identify inactive identities

You can create a report in the IBM Cloud console to identify which users, service IDs, trusted profiles, and API keys in your account are inactive. Removing access for inactive identities can reduce the risk of unauthorized access to your IBM Cloud® resources and help you manage access more efficiently. For more information, see Identifying inactive identities.

Updated process for assigning access

Assigning IAM, Classic Infrastructure, and IBM Cloud Foundry access just got more streamlined. When assigning access, each service that you select has an in-context description. You can also find all IAM access policy and access group information for an identity under a single tab. Check out the updated process by assigning access to any user, service ID, or trusted profile.

Assign a policy for All IAM Account Management services

You can now assign an access policy for All IAM Account Management services, which includes the IAM Identity service, IAM Access Management service, IAM User Management service, IAM Access Groups service, and future IAM services. By assigning an access policy for a group of services, you decrease the number of policies in your account and reduce the time and effort to manage access.

13 June 2022

Updated process for onboarding services in Partner Center

Partner Center has made it easier to provide required support information. The process has improved the approval process for publishing your product, and it has added consistency to all third-party product information which helps customers find the help that they need.

For more information, see Defining your support experience.

May 2022

04 May 2022

Required access to view service credentials

When the credential level access can't be determined by comparing the access of the user and the credential, the credential is redacted. Services that support access control on resources that are allocated in the context of a service instance, such as Cloud Object Storage buckets, now require the user to have the resource-controller.credential.retrieve_all action to view service credentials.

To determine if your service is affected, review your service's documentation. For more information about the required access to view service credentials, see Viewing a credential.

April 2022

5 April 2022

Onboarding virtual server images for Power Virtual Server

You can now onboard virtual server images for Power Virtual Server to private catalogs and the IBM Cloud catalog. For more information about onboarding to the IBM Cloud, see Registering a virtual server image for Power Systems in IBM Cloud Partner Center. For more information about onboarding to your private catalog, see Onboarding a virtual server image for Power Systems to a private catalog.

Upload translations of your software by using the CLI

You can now download and upload translations of your software by using the CLI. For more information, see Translating product details by using the CLI.

4 April 2022

Adding custom service parameters for your service

If the provisioning process of your product requires additional information from your customers, you can now add custom input fields for your product in IBM Cloud Partner Center. For more information, see Adding custom service parameters for your service in Partner Center.

March 2022

25 March 2022

Trusted profiles are now members of access groups

You can now add trusted profiles as members of access groups like other IAM identities, such as users and service IDs. For more information, see Creating trusted profiles and Setting up access groups.

1 March 2022

Use feature and subscription codes to create new accounts

When you register for a new IBM Cloud account, you can use a feature code that you received from participating in a special event or a subscription code that you received by email. Instead of verifying your identity by entering your credit card information, click the option to **Register with a code** to complete the set up of your new account.

February 2022

28 February 2022

365 days of status history is now available

You can now easily view events that were completed in the past year on the IBM Cloud Status page. Previously, you could view only the history of completed events from the past month. For more information about the IBM Cloud Status page, see Viewing cloud status.

7 February 2022

Support for onboarding third-party services in Partner Center

Onboarding new services in the resource management console is no longer supported. You must use Partner Center to onboard new services to the IBM Cloud console. You can continue to manage existing services in the resource management console.

December 2021

17 December 2021

Restricting domains for account invitations

You can now restrict membership to your account based on the domain of the users by creating an allowlist. This way, only users with a specific domain or domains can be invited to the account. For more information, see Restricting domains for account invitations.

November 2021

8 November 2021

Data centers closed in 2021

The following list shows the locations with the associated data centers and specific PODs that were closed by 08 November 2021.

- Dallas: DAL05 – POD2
- Houston: HOU02
- Melbourne: MEL01
- Oslo: OSL01
- Washington DC: WDC01 – POD3 and POD4

For more information about upcoming data center closures, see Data center consolidations.

1 November 2021

Accessing Partner Center from the IBM Cloud catalog

As a third-party provider, you can now click **Sell on IBM Cloud** to access Partner Center directly from the IBM Cloud catalog. For more information, see the catalog.

October 2021

25 October 2021

Securing IBM Cloud accounts

Our focus is to provide the most secure public cloud. That's why we will be verifying user identities and securing accounts through credit card verification when creating new accounts starting 25 October 2021. Don't worry, we won't charge you for signing up, and you can still try IBM Cloud for free. For more information, see Securing IBM Cloud accounts.

05 October 2021

Suspending and deprecating private products

As a private catalog owner, you can now suspend or deprecate software that's in your private catalogs. When you suspend a version of software, it's immediately removed from your private catalog. When you deprecate software, it remains available to users with access to the private catalog for 90 days. After 90 days, it's permanently removed.

For more information, see Deprecating a private product and Suspending a version of a private product.

Suspending and deprecating third-party software

As a third-party provider, you can now suspend and deprecate published software from the IBM Cloud catalog to meet the needs of your product's lifecycle. When you suspend a product or version, it is immediately removed from the catalog. When you deprecate a product or version, it remains available for use in the catalog for 90 days. After 90 days, it's permanently removed.

For more information, see Deprecating software from the IBM Cloud catalog and Suspending your product from the IBM Cloud catalog.

Filtering the IBM Cloud catalog by specific providers

The latest enhancements to IBM Cloud catalog include support for filtering products by provider name. If you're looking for a specific provider's products or curious about how many products that a provider offers, you can use the **Provider** filter to narrow down your search. To explore the filter updates, see the catalog.

September 2021

27 September 2021

Securing your resources with context-based restrictions

You can now use context-based restrictions to configure and enforce access restrictions for IBM Cloud resources. The access restrictions are based on the network location where an access request is created. These restrictions work in tandem with traditional IAM access policies, which are based on identity, to provide an extra layer of protection. Since both IAM access and context-based restrictions must enforce access, context-based restrictions offer protection even in the face of compromised or mismanaged credentials.

You can create context-based restrictions by defining one or more network zones, which contain allowed network locations, and then associating the zones with the cloud resource through a context-based restrictions rule. Network zones can be defined in terms of IP address constructs, VPCs, and service references, which grant access to a resource from another cloud service.

For more information, see What are content-based restrictions?

24 September 2021

Assigning access policies based on resource location

For supporting services, like Container Registry, you can now scope access to more granular locations in your policies, such as geography, country, metro, satellite location.

20 September 2021

Checking the root cause of an incident

You can now visit the IBM Cloud Incident reports page to check the health of an IBM Cloud service. You can find Customer Incident Reports (CIR) on the page, which provide Root Cause Analysis (RCA) for past major outages. You are able to view and download the incident reports from the page about any events that affect IBM Cloud availability. These reports are in PDF file formats and available for 5 years starting from the date when the event happened.

For more information, see Checking Incident reports.

01 September 2021

Update to the latest CLI version

The IBM Cloud CLI team is deprecating the current CLI plugin repo infrastructure for downloading and updating the CLI binary files and plugins. The CLI is migrating to a new infrastructure that works with IBM Cloud CLI version 2.0.0 or newer. Ensure that you update all instances of the IBM Cloud CLI before 1 October 2021 to avoid disruptions for downloading or updating the plugins. No impact is expected for users who don't need to update any plugins or the CLI.

August 2021

05 August 2021

Delivering notifications by using Microsoft Teams webhooks

Adding Microsoft Teams webhooks to your distribution list is available for receiving account-related IBM Cloud notifications. To create a webhook, you need to add an incoming webhook to a Teams channel and a unique URL that maps to the selected channel. With this webhook integration, you can easily receive the notifications in a selected Microsoft Teams channel in which you added the incoming webhook.

For more information, see Adding Microsoft Teams webhooks to a distribution list .

04 August 2021

Support for third-party Operator bundles from GitHub

Third-party providers can now onboard Operator bundles from GitHub repositories to IBM Cloud by using TGZ files. Previously, the ability to onboard Operator bundles was limited to Operator bundles in Red Hat OpenShift on IBM Cloud registries. This capability is supported by the beta release of IBM Cloud Partner Center and our general availability release of private catalog capabilities.

Software onboarding enhancements: progress indicators

When you onboard software to the IBM Cloud catalog or private catalogs in your account, you'll find new indicators in the console that you can use to track your progress with configuring and validating your product. For more information, see Onboarding software to your account.

03 August 2021

Filtering the IBM Cloud catalog by product type

When you navigate to the IBM Cloud catalog, your view by default includes all types of products: services, software, and professional services. To help you quickly find the product that you're looking for, you can now filter the products to view services only, software only, or professional services only. For more information, see IBM Cloud catalog.

July 2021

27 July 2021

Assigning access to federated users and compute resources by using trusted profiles

You can use trusted profiles to automatically grant federated users in your account access to resources with conditions based on SAML attributes from your corporate directory. You can also use trusted profiles to manage the authorization of applications that are running in compute resources, such as IBM Cloud Kubernetes Service, to access other IBM Cloud services without the need for service IDs or API keys. For more information, see

Creating trusted profiles.

15 July 2021

Scoping support cases to what matters to you

When you create a case in the IBM Cloud Support Center, you now have options to narrow the subject of your case to a specific topic that's most closely related to the issue you're experiencing. As a result, you can ensure that your support case gets routed to the appropriate support engineer and resolved as efficiently as possible. For more information, see Creating support cases.

14 July 2021

Delivering notifications by using Slack webhooks

In addition to generic webhooks, you can now add Slack webhooks to your distribution list and receive account-wide IBM Cloud notifications through them. To create a webhook, you need to set up an app in Slack and a unique URL. With this new webhook integration, you can easily receive the notifications in a selected Slack channel in which you installed your app.

For more information on Slack webhooks, see Adding Slack webhooks to a distribution list.

07 July 2021

Service-level discounts for accounts with the Pay as you go with Committed Use billing model

If you have an account with the Pay as you go with Committed Use billing model, you might be eligible to receive service-level discounts. You can apply service-level discounts to commitments in addition to the platform-level discount.

Service-level discounts are provided by IBM Cloud Sales to customers with the Pay as you go with Committed Use billing model.

Support for third-party Operator bundles from Red Hat OpenShift on IBM Cloud registries

Third-party providers can now onboard Operator bundles from Red Hat OpenShift on IBM Cloud registries to IBM Cloud. The Red Hat OpenShift on IBM Cloud registries include Certified Operators, Marketplace Operators, and community Operators. This capability is supported by the beta release of IBM Cloud Partner Center and our general availability release of private catalog capabilities. For more information, see Onboarding a Certified Operator bundle from a Red Hat registry.

June 2021

01 June 2021

New invitation flow for existing IBM Cloud users

To enhance security and user protection, IBM Cloud® now requires all users to accept an invitation in order to become an active user within a new account. The new invitation flow has an impact only on inviting existing IBM Cloud® users. Previously, existing users were being automatically onboarded to each new account as they were invited. After this change, these users need to accept an invitation in their notifications, by email, or by using the CLI to onboard to a new account.

To accept invitations in the CLI, existing members of IBM Cloud® must use the `ibmcloud login` command. They need to target the account that they are invited to join and use the new `--accept` flag.

Note: As an account administrator, you might want to remind your users to accept these invitations.

May 2021

07 May 2021

Support for third-party virtual server images with Terraform

Third-party providers can now offer virtual server images deployed by using Terraform in IBM Cloud. This capability is supported by the beta release of IBM Cloud Partner Center and our previously general availability release of private catalog capabilities. For more information, see Onboarding a virtual server image with Terraform.

April 2021

21 April 2021

Delivering notifications by using webhooks

You can now easily add webhooks to the notification distribution list in addition to adding email addresses. You can register a webhook with your account on the Notification distribution list page to receive all account notifications. Administrators can use webhooks to configure an application to receive asynchronous notifications whenever an event occurs on the IBM Cloud platform. Any registered webhook must support HTTP POST requests and accept the notification as a JSON.

For more information, see Adding webhooks to a distribution list.

19 April 2021

Upcoming changes to the user invitation flow

Starting June 01 2021, to enhance security and user protection, all users will be required to accept an invitation to become an active user within a new account. The new invitation flow will have an impact only on inviting existing IBM Cloud® users. Existing users are currently being automatically onboarded to each new account as they are invited. After this change, these users will need to accept an invitation in their notifications, by email, or by using the CLI to onboard to a new account.

Concerned about how this change will impact your automation? To avoid any disruption to on-going workflows, you need to check your scripts:

- To accept invitations in the CLI, existing members of IBM Cloud® must use the `ibmcloud login` command. They need to target the account that they are invited to join and use the new `--accept` flag.

Note: As an account administrator, you may want to remind your users to accept these invitations upon initial change of this behavior.

March 2021

31 March 2021

Support for third-party Operators deployed on Red Hat OpenShift on IBM Cloud

Third-party providers can now offer Operators in IBM Cloud. This capability is supported by the beta release of IBM Cloud Partner Center and our previously general availability release of private catalog capabilities. For more information, see Onboarding a Node-Red Operator.

New catalog filter for Financial Services Validated services

You can now search the IBM Cloud catalog for services that are designated as Financial Services Validated. Services with this designation leverage the industry's highest levels of encryption certification, provide preventive and compensatory controls for financial services regulatory workloads, multi-architecture support and proactive, and automated security. See IBM Cloud Framework for Financial Services for more information.

26 March 2021

IBM Cloud Code Engine is supported as an app deployment type

IBM Cloud® Code Engine is now supported as an application deployment type when you use a starter kit in the IBM Cloud console.

Code Engine is a fully managed, serverless platform that runs your containerized workloads, including web apps, micro-services, event-driven functions, or batch jobs. Code Engine even builds container images for you from your source code. Because these workloads are all hosted within the same Kubernetes infrastructure, all of them can seamlessly work together. The Code Engine experience is designed so that you can focus on writing code and not on the infrastructure that is needed to host it.

Duplicate access group names

Access groups must use unique names. You can't create or update IAM access groups to use the same name.

February 2021

26 February 2021

Managing user login sessions

You can customize the duration of working sessions for user's on your account. Customize the duration of users active settings and select the amount of time a user can be inactive before they are signed out and need to enter their credentials again. Update your accounts Identity and Access (IAM) log in sessions from Settings page. For more information, see Managing user's log in session durations.

25 February 2021

Controlling access to resources by using tags

You can now create tags to control access to your resources and your team's projects can grow without requiring updates to your IAM policies. Incorporating tags into access policies gives you the ability to share a resource across multiple projects and you can change access by simply changing the tags on a resource. For more information, see Controlling access to resources by using tags.

08 February 2021

New notifications experience

A more detailed Notification preferences page is now available for you to customize your preferences for receiving email notifications. You receive only one email per event unless you subscribe to them, or you can subscribe to specific incidents from the Status page on an ad hoc basis. For more information, see Setting email preferences for notifications.

Based on which preferences the account owner or administrator sets, users in the account can view all IBM Cloud incidents, maintenance, announcements, and security bulletins on the Notifications page. They can filter the list by selecting a specific type of event, or by using keyword searches. For more information, see Viewing notifications.

01 February 2021

Managing product availability in catalogs by location

As the account owner or administrator, you can manage access to products in both the IBM Cloud catalog and the private catalogs in your account based on the location in which the products are deployed. For instance, if you want to limit access to products that are deployed in the Dallas 1 (us-south-1) zone, you can set a filter to include only those products. For more information, see Managing catalog settings.

January 2021

11 January 2021

Pay as you go with Committed Use pricing model

Customers with a Subscription account can use the new pricing model, IBM Cloud Pay as you go with Committed Use. The new pricing model provides you with additional benefits as you navigate and build on IBM Cloud.

With this pricing model, you commit to spend a certain amount and receive discounts across the entire platform. You are billed monthly based on your usage, and unlike a subscription, you continue to receive a discount even after you reach your committed amount. For more information, see Pay as you go with Committed Use pricing model .

December 2020

18 December 2020

Managing features and other updates to IBM Cloud Shell

Account owners or users with Cloud Shell administrator access can enable or disable Cloud Shell features for an account. The available features in this release are **File upload and download** and **Web preview**. The feature settings apply only to the enabled Cloud Shell locations. For more information, see Enabling or disabling Cloud Shell features for an account .

An account administrator can grant specific users access to Cloud Shell and its features, even if Cloud Shell settings are disabled at the account level. For more information, see Assigning access to Cloud Shell and its features at a user level .

The following new service roles are available:

- Cloud Operator
- Cloud Developer
- File Manager

For more information, see IAM roles and actions .

IBM Cloud Shell now uses a Red Hat™ Linux™ bash shell instead of a x86-64 Ubuntu Linux™ bash shell.

For a complete list of changes, see the Cloud Shell release notes .

17 December 2020

IBM Cloud CLI Version 1.3.0 is now available

The 1.3.0 release of the IBM Cloud CLI includes:

- Initial support of private endpoints in two regions: `us-east` and `us-south`.
- Anonymous usage statistics are no longer collected for those that had previously opted in to usage statistics collection. The `--usage-stats-collect` flag is removed from the `config` command.
- Improved performance of several commands involving resource groups.
- New command to toggle the Intelligent Platform Management Interface (IPMI) on and off: `sl hardware toggle-ipmi`.
- The add `--output` flag is added to the `resource service-key-create` command.
- For the `iam access-groups` command, `Public Access` is included in the output.
- Upgrade to Go language `Golang 1.15.5`.

For more details about Version 1.3.0, see the IBM Cloud CLI release notes.

For more information about the CLI, see Getting started with the IBM Cloud CLI.

10 December 2020

Enhanced payments & invoicing for new US-based Pay-As-You-Go accounts with credit card billing

We are excited to announce our latest unification project for new US-based IBM Cloud Pay-As-You-Go accounts with credit card billing. For increased clarification and consistency, the unified experience includes the following enhancements:

- A single invoice that includes comprehensive usage details for all your products
- The ability to download the invoice directly from the console
- The ability to update a credit card in the console without being redirected to a different website
- A one-to-one mapping between your invoice and your usage dashboard in the console

For more details, including a video walk-through of the enhancements, see The Enhanced Unified Billing and Payment Experience in IBM Cloud.

04 December 2020

Enhancements to the IBM Cloud Status page

You can now experience an enhanced version of the Status page in the IBM Cloud console that offers reduced time to visibility of issues, as well as more detailed and up-to-date status information. For more information, see Viewing cloud status.

November 2020

30 November 2020

Data centers closed in 2020

The following list shows the locations with the associated data centers and specific PODs that were closed by 30 November 2020.

- Dallas: DAL07
- Seattle: SEA01
- Legacy Planet: D2 (colocation in Dallas), D6 and D7 (Dallas), H2 (Houston)

For more information about upcoming data center closures, see Data center consolidations.

25 November 2020

Support for U2F MFA and other MFA factors

By default, users in your account authenticate themselves by logging in with a username and password. To require users to use more secure authentication factors, the following MFA options are now available.

- MFA for users with an IBMid: Users authenticate by using an IBMid, password, and time-based one-time passcode (TOTP). This option applies to all users or just non-federated users.
- MFA for all users: This option applies to users who are using either an IBMid or an external identity provider (IdP). Users authenticate by using email-based MFA, TOTP MFA, or U2F MFA. The U2F MFA factor is based on the FIDO U2F standard, and it offers the highest level of security.

For more details, see Enabling MFA for your account.

October 2020

28 October 2020

Enhanced Support Center

The Support Center is now updated to help improve your experience with creating and managing your support cases. You can refine the scope of your cases by routing them to a specfic resource and provide feedback on your cases. You will also find popular FAQs that are featured based on your specific issue. For more details, see Announcing the Release of Our New Support Center Enhancements .

September 2020

18 September 2020

Managing IBM Cloud Shell settings and other updates to IBM Cloud Shell

Account owners or users with Cloud Shell administrator access can manage Cloud Shell settings from the IBM Cloud console. For more information, see Updating Cloud Shell settings.

For a complete list of changes, see the Cloud Shell release notes.

Restricting account access by using IAM account settings

For increased control over which users can access your account and work with API keys and service IDs, leverage the three new settings that are available on the **Manage** > **Access (IAM)** > **Settings** page in the console.

- Restrict access to your account to only users coming from a specified IP address or range that you set. For more information, see Allowing specific IP addresses.
- Block all users from creating API keys in the account except for those that you give explicit access. For more information, see Restricting users from creating API keys.
- Block all users from creating service IDs in the account except for those that you give explicit access. For more information, see Restricting users from creating service IDs.

August 2020

31 August 2020

Introducing Security and Compliance Center

For highly regulated industries, such as financial or healthcare services, achieving a continuously secure and compliant cloud environment is an important first step toward protecting customer and application data. But, we understand that historically this has been a difficult and manual process, potentially placing your organization at risk, which is why, we're excited to introduce the Security and Compliance Center - built directly into the IBM Cloud platform.

With the Security and Compliance Center, you can embed checkpoints into your everyday workflows hat not only help to monitor for security and compliance but enforce configuration standardization across our accounts. By ensuring you have automation in place that monitors for risk, you can identify security vulnerabilities quickly and work to mitigate the impact that they have to your business.

For more information, see Getting started with Security and Compliance Center.

20 August 2020

IBM Cloud CLI Version 1.2.0 is now available

The 1.2.0 release of the IBM Cloud CLI includes:

- New commands under `sl loadbal` namespace to support the IBM Cloud® Load Balancer service.
- Output CSV support for billing commands.
- Validation of service instance ID input in policy commands under `iam` namespace.
- Additional deployment capabilities with the `ibmcloud dev` command, including manual deployment to Knative and deployment to Red Hat® OpenShift® on IBM Cloud® containers.
- Deprecation of the `cfee` namespace.

For more information about the IBM Cloud CLI, see Getting started with the IBM Cloud CLI.

12 August 2020

Increased usage quota and other updates to IBM Cloud Shell

With this update to IBM Cloud Shell, the weekly usage quota was increased from 30 hours a week to 50 hours a week so that you can use Cloud Shell even more. Also, Cloud Shell was changed to have separate workspaces for each user and account, whereas previously each user's workspace was shared across all of their accounts. In your sessions, you'll also now have access to more command-line tools with the addition of the Operator SDK for Kubernetes, the Mercurial source content management tool, and the Bazaar version control system, as well as updates to existing tools.

For a complete list of changes, see the Cloud Shell release notes.

10 August 2020

Customizing scoped dashboards

Create and manage customized dashboards that can be used to offer an organized and relevant workspace. You can share the dashboards with different users in your account or as a way to group resources for specific projects or teams. When you create a new dashboard, you can start from scratch with a blank template or choose one of the other available templates. For more information, see Working with scoped dashboards.

04 August 2020

Audit logs for private catalogs

View recorded changes to your private catalogs with the new audit logs feature. Audit logs are recorded when you create, delete, or update your private catalog. Visit your Audit logs page in the IBM Cloud console by going to **Manage** > **Catalogs**, and selecting **Audit logs**. For more information, see Viewing changes to your private catalogs.

July 2020

07 July 2020

Notification distribution list

You can now create a list of up to 10 email addresses that receive account-wide notifications in the IBM Cloud console by going to **Account** > **Notification distribution list**. Users that are added to the distribution list are notified about any event that's affecting the account. For more information, see Adding users to a distribution list.

CLI support for catalog filtering in IBM Cloud enterprises

You can now use the following commands to set and manage filters in accounts within an enterprise hierarchy:

- `ibmcloud catalog filter get`

- `ibmcloud catalog filter create`
- `ibmcloud catalog filter offering`
- `ibmcloud catalog filter delete`

Each command operates the same at the enterprise level by default. You can also apply the commands to specific account groups in an enterprise. For more information, see the Catalogs management CLI plug-in.

Professional services are available in the IBM Cloud catalog

We now offer professional services that aim to help guide you in transforming your business. This strategy combines expertise, tools, and methodologies to develop cloud strategies and implementation plans that align with your business goals.

To explore what's available, go to the catalog, and click **Professional services**.

Catalog management SDKs

Four new Software Development Kits (SDKs) for the catalog management service, which includes private catalogs and software products, are now available.

- The IBM Cloud Platform Services Java SDK
- The IBM Cloud Platform Services Node.js SDK
- The IBM Cloud Platform Services Python SDK
- The IBM Cloud Platform Services Go SDK

For more information, see the Catalog Management API reference.

OVA images are available in the IBM Cloud catalog

With the deployment power of the IBM Cloud Schematics service, support for third-party OVA images targeting VMware vCenter Server deployments is now available. To quickly find the available products in the catalog, click **Software**, and select **OVA Images** from the **Software** list.

June 2020

24 June 2020

IBM Cloud Shell is generally available

IBM Cloud Shell, the browser-based shell environment, is now generally available (GA)! This release includes greater region support with the addition of the Tokyo (`jp-tok`) region. Cloud Shell sessions now also support inputting and viewing double-byte characters on the command line, so you can work in languages such as Japanese, Simplified and Traditional Chinese, and more. The server image was also updated to include updates to several command line tools, in addition to the regular updates for the IBM Cloud CLI and plug-ins.

To try out IBM Cloud Shell, see Getting started with IBM Cloud Shell.

May 2020

21 May 2020

Connect to an external identity provider for authentication

You can now connect your external identity provider to an IBM Cloud App ID instance, and then configure that App ID to connect directly to IBM Cloud Identity and Access Management (IAM) to federate authentication for users in your enterprise to your IBM Cloud account.

By using this new integration between App ID and IBM Cloud IAM, you can simplify the log in experience for users in your enterprise and automatically add users to your account through their login, instead of having to invite or federate each user individually with IBMid. For more information, see Enabling authentication from an external identity provider.

01 May 2020

Support for catalog filtering in IBM Cloud enterprises

You can now use filters to customize which products in the IBM Cloud catalog are available in accounts within an enterprise hierarchy. Filters that are set at a parent account level apply to all child account groups and accounts. For more information, see Managing products for an IBM Cloud enterprise.

April 2020

21 April 2020

Unified notifications experience

Viewing your notifications is now easier than ever with the new unified notifications experience. The notifications page is the centralized place to view and manage all incidents, maintenance, and announcements that might affect your account. In the IBM Cloud console, click the **Notifications** icon on the console menu bar to view your notifications.

Additionally, you can find all infrastructure notifications on the new notifications page. To learn more about the new notifications experience, see Viewing notifications.

18 April 2020

New features and multi-region support for IBM Cloud Shell (Beta)

This update to IBM Cloud Shell (Beta) includes several new features to help you work from the command line in IBM Cloud:

- You can now preview web apps within Cloud Shell. By running apps on an available port, you can open a preview at a URL that only you can see, similar to a cloud version of `localhost`.
- You can now use Cloud Shell within the Frankfurt region in addition to the existing Dallas region.
- Usage timeouts were changed to extend how long you can leave your sessions idle before they're closed. And, there's no longer a time limit on continuous usage.
- An update to the server image adds the GNU Automake (`automake`) and GNU Compiler Collection (`gcc` , `gcov` , and `gcov-tool`) tools to your sessions.

For a complete list of changes, see the Cloud Shell release notes.

March 2020

31 March 2020

IBM Cloud CLI Version 1.0.0 is now available

The 1.0.0 release of the IBM Cloud CLI includes two new global options to help you more easily automate command-line scripting tasks. The new `--output json` option provides command output in a parsable JSON format, and the `--quiet` option suppresses some extra messages such as progress indicators simplify automated processing. The IBM Cloud CLI now also includes the IBM Cloud Developer Tools (`ibmcloud dev`) commands, so you can work with apps, pipelines, toolchains, and more without needing to install a separate plug-in. The Cloud Foundry CLI (`ibmcloud cf`) is no longer bundled but can still be installed separately.

For more information about the IBM Cloud CLI, see Getting started with the IBM Cloud CLI and Developer Tools.

18 March 2020

Customizing how access is assigned by using custom roles

Until now, you assigned access by using the available platform and service roles that have specific actions that are mapped by individual services. Sometimes, you must assign multiple roles to achieve the required access for your users. With the latest IBM Cloud Identity and Access Management feature, you can now customize how access is assigned within individual services by creating custom roles. You can choose to organize any number of actions available for a specific service into a new role with a name of your choosing. For more information, see Creating custom roles.

February 2020

26 February 2020

Updated handling of API keys for removed users

When you remove a user from your account, IBM Cloud automatically cleans up the API keys associated with that user's identity, so you don't have to. For more information, see the Updated handling of IBM Cloud API keys for users removed from accounts blog post.

January 2020

21 January 2020

IBM Cloud Shell (Beta) is now available

We're introducing a new way to access IBM Cloud from the command line with no installation needed - IBM® Cloud Shell! IBM Cloud Shell is a cloud-based shell environment that gives you instant command-line access through your web browser. It's preconfigured with the full IBM Cloud CLI, all CLI plug-ins, and tons of tools so you have everything that you need to manage apps, resources, and infrastructure. To try out this beta release, see Getting started with IBM Cloud Shell.

09 January 2020

Improved Support Center

The latest enhancement to the support center offers a personalized experience to better resolve any IBM Cloud related issue. Navigate through the landing page to view incidents that are specific to your account along with a list of your open cases. Additionally, our self-help options have expanded, popular FAQs and recommended topics are populated to provide information that is relevant to your account. To check out the new experience, log in and go to the Support Center .

December 2019

19 December 2019

Disable public access to resources

All users and service IDs are members of the Public Access group by default, which provides unauthenticated access to the resources that the group can access. You can now disable public access at the account level for cases in which you might want to prevent policies that could publicly expose resources from being created. For more information, see Managing public access to resources.

November 2019

26 November 2019

Delete resource groups

Besides the default resource group that's automatically created for you when create your IBM Cloud account, you can now delete any resource group that doesn't contain any resources. For more information, see Managing resource groups.

04 November 2019

Bitnami Helm charts

IBM and VMware are partnering to provide the Bitnami Application Catalog to IBM Cloud customers. The first set of offerings available in the IBM Cloud catalog are the Bitnami Helm charts, which Bitnami has created and validated to work on the IBM Cloud Kubernetes Service. Kubernetes developers typically install Helm charts manually from the CLI by using the helm command. However, this installation process has been simplified by building an auto-installation feature in the IBM Cloud catalog. Check out these offerings today by selecting **Bitnami** from the **Provider** filter.

October 2019

18 October 2019

Apply promo codes to your account and services

IBM Cloud now has promo codes that you can use to get credit toward your account and services. Promo codes are provided by IBM Cloud sales on a limited basis to customers with billable accounts. You can apply promo codes from a new Promotions page or from the catalog when you create a resource in the IBM Cloud console. For more information, see Applying promo codes.

Download access reports for specific resources

Have you ever wanted to get a quick view of all users, service IDs, access groups, and services that have access to a specific resource in your account? Well, now you can download a report that includes details about who can access your resource. From the My resources page in your account, you can select the **Export access report** option in the row for an IAM-enabled resource to get the report. For more information about the types of data available in the report and who has access to view the report, see Exporting an access report.

15 October 2019

Mapping actions to IAM roles

When you invite a user to your account or assign an existing user IAM access, you can review each action that is mapped to a role. Click the **Actions for role** option to view a list of all actions that are mapped to a specific role when you're selecting the roles to assign in an access policy. By reviewing the action to role mappings, you can have confidence that you are always assigning the correct level of access to users in your account. For more information about this new enhancement, see Invite User Flow and Transparent Actions.

03 October 2019

Cloud Paks and Schematics templates

With the deployment power of the IBM Cloud Schematics service, we now deliver bundled solutions through packaged software and deployable automation scripts that empower you to build, manage, and modernize your cloud architectures faster and more securely. The first release of these offerings include IBM Cloud Paks and a handful of useful Terraform-based templates.

From the catalog, filter by the offering type and deployment target to find what you're looking for faster. Select **Cloud Paks** and **Terraform** to check out the new offerings that are available.

September 2019

23 September 2019

Term-based model for IBM Cloud support subscriptions

New subscriptions for IBM Cloud support now follow the same term-based model as subscriptions for IBM Cloud platform services. Rather than credit being valid in monthly increments, credit is now available up front for an entire subscription term, which can be up to one year's worth of credit. You now have more flexibility to use your support credit when you need it, and the chances of incurring monthly overages are reduced. When you buy or renew a support subscription, any existing active support subscriptions are converted to this new flexible model for the remainder of their term. For more information, see Support subscriptions.

You can also now view your support subscriptions in the IBM Cloud console so you can keep track of your available credit. For more information, see Viewing your support costs.

12 September 2019

Redirecting SoftLayer to IBM Cloud

SoftLayer account owners who previously didn't have access to the IBM Cloud platform can now manage their infrastructure, services, and applications from one location: cloud.ibm.com.

July 2019

25 July 2019

IBM Cloud enterprises for centrally managing multiple accounts

You can now centrally manage billing and usage for multiple accounts by creating an IBM Cloud enterprise. With an enterprise, you can create a multitiered hierarchy of accounts by organizing related accounts into account groups. Enterprises simplify management of multiple accounts with the following key features:

- Consolidated billing means that you can manage billing, invoicing, and payment for all accounts from a single place, the enterprise account.
- Subscription credit is aggregated into a credit pool and shared with all accounts in the enterprise. Not only is tracking your subscriptions easier, but you can get fewer, larger subscriptions for a better discount because the credit is shared.
- Top-down usage reporting gives you a unified view of usage costs from all accounts, organized according to your enterprise hierarchy.

If you have multiple accounts, at least one of which is a Subscription account, you can create an enterprise. See What is an enterprise? and Introducing IBM Cloud Enterprises for more information.

Subscriptions page for tracking subscription credit spending

If you have a Subscription account, you can now view all of your subscriptions and analyze your credit spending on the Subscriptions page. You get a high-level view of the total subscription credit in your account and detailed charts that visualize trends such as your credit burndown and monthly spending. You can also view credit from any promotions in your account. For more information, see Managing subscriptions.

Additionally, to better reflect their usage, codes that you apply to add subscription credit to your account are now called subscription codes rather than feature codes.

02 July 2019

Managing SoftLayer SAML federation on IBM Cloud

Former SoftLayer users who set up a SAML identity provider for logging in with federated IDs can now manage their configuration data in the IBM Cloud console in Access (IAM) on the Identity providers page. This type of federation is deprecated, so new identity providers can't be set up

currently. You can continue to update your identity provider data or choose to delete this federation in favor of switching to federating with IBMid.

June 2019

14 June 2019

Customize your dashboard

You can now control what's displayed on your dashboard. Customizing your dashboard includes the ability to add, remove, and rearrange the widgets. For more information, see Customizing your dashboard.

April 2019

04 April 2019

Export usage data with associated tags

You can now use our newest tagging capabilities to manage resources, usage, and costs in the exported usage report. When you add a tag to a resource, you can view the tag that is associated with the resource. In the IBM Cloud console, go to **Manage**> **Billing and Usage**> **Usage**> **Export CSV**> **Instances** to download your usage report.

March 2019

25 March 2019

Enabling public access to resources

You can now enable public access to objects in your IBM Cloud® Object Storage buckets by using a new access group that is provided for you in your account. This new access group is called the `Public access` group, and all users and service IDs are added to it by default. You can update the policies for the access group to enable all users, even unauthenticated users, access to the resource that you specify in the policy. Learn more about the public access group.

12 March 2019

Multifactor authentication for users with federated IDs

Account owners or users assigned the administrator role for the billing account management service can enable multifcator authentication (MFA) for all users in their account. Federated users who use their corporate or enterprise single sign-on ID can now be required to authenticate by using MFA for logging in to IBM Cloud. For more information about this feature enhancement and what you need to know about enabling MFA for your account, see Introducing MFA for IBM Cloud Users with Federated ID .

December 2018

31 December 2018

New appdomain.cloud host name option

A new host name option `*.appdomain.cloud` is available on cloud.ibm.com.

Previously, the `mybluemix.net` domain was used for hosting apps in various deployment targets, such as IBM Cloud Kubernetes Service or Cloud Foundry. Any apps that are hosted on `mybluemix.net` are not impacted.

The subdomain for Cloud Foundry apps is `cf.appdomain.cloud`. The subdomain for apps that you deploy to IBM Cloud Kubernetes Service is `containers.appdomain.cloud`.

November 2018

30 November 2018

New Cloud Foundry API endpoints

The legacy `api.*.bluemix.net` Cloud Foundry API endpoints are still available for compatibility with earlier versions. However, you can update scripts and infrastructure automation to use the following new Cloud Foundry API endpoints for your region:

- api.us-south.cf.cloud.ibm.com (previously api.ng.bluemix.net)
- api.eu-gb.cf.cloud.ibm.com (previously api.eu-gb.bluemix.net)
- api.us-east.cf.cloud.ibm.com (previously api.us-east.bluemix.net)
- api.eu-de.cf.cloud.ibm.com (previously api.eu-de.bluemix.net)
- api.au-syd.cf.cloud.ibm.com (previously api.au-syd.bluemix.net)

Enhanced support experience for IBM Cloud

With the Support Center, you can work to resolve all IBM Cloud-related issues. The landing page provides you FAQs, so you can find the answer to your question without even contacting IBM Cloud. You can also chat with a live support rep. Your cases can now be managed in from a single location. In the console, go to **Support** > **Manage cases** to create, view, or edit cases.

You can also find the Status page from the Support Center. It is enhanced to include all unplanned incidents, planned maintenance, announcements, and security bulletin notifications about key events that affect the IBM Cloud platform, infrastructure, and major services. Click **View cloud status** from the Support Center. To check out the new experience, log in and go to the Support Center .

Unified login, API keys, and user and access management

With our latest updates, you can take advantage of a simplified secure login that is available for all users regardless of your ID type. Whether you have an IBMid or a SoftLayer ID, you can quickly log in to the IBM Cloud console from our enhanced login page. You can also make secure API calls across IBM Cloud and automate your CLI login by using an IAM API key or an IAM access token.

After you log in, you can now see all users, including platform and classic infrastructure users, from your Users page in the Access (IAM) UI. Depending on your access to view other users in the account, you can filter your view quickly by account users, classic infrastructure users, or Cloud Foundry org. You can also use the filters to find users quickly by name, email, or status.

Now that all of your users are in a single console, you can manage their access to all types of resources from the same place. Access starts with the user, so start by selecting a user from your list. Then, depending on which type of resource that you want to assign access to, you can choose from IAM access policies, Cloud Foundry access, or classic infrastructure permissions. If you want to assign IAM access policies, try creating an access group to streamline your access management process by adding all users to the same access group that need the same policies assigned.

For more details, check out Outstanding User Access Improvements Help Deliver a Unified IBM Cloud Platform .

Find all IBM Cloud CLI plug-in documentation in one place

You can now access all of the IBM Cloud CLI plug-in documentation in one location, making it easier for you to find any CLI command that you are looking for. Check out the References section in the CLI documentation.

New dashboard and Resource list page

With our latest update, you can now view all your platform and infrastructure services from one location. When you log in, you can check out the new dashboard. After you add resources to your account from the catalog, you can use your resource list to get a full view of your account resources:

- The dashboard is redesigned so that you can view a summary your resources, maintenance, status, apps, support, usage, and users.
- You can find more details about your resources by going to your resource list. You can tag your resources to organize them, or select them to update the details page.
- Now that you can see all of your resources in one place, we added a global search so that you can quickly find resources that you created and expect to find in your resource list.

- You can also search for catalog results, so you can quickly find resources to add to your account.

Unified account, billing, and user profile information

Your account, billing, and profile information is now simplified. You can view your account information for all of your platform and infrastructure resources in a unified console.

Your profile and settings area contains information about you as well as your email notification preferences for all resource types.

Your account information area contains information about your company or organization, account settings, and quick access for working with resource groups and Cloud Foundry orgs. You can even find best practices to help you get up and running quickly!

Your billing and usage area of your account helps you understand your bill, make payments, monitor subscriptions, get quotes, track orders, and set spending notifications.

Check out Bringing It All Together: A Single Account and Billing Management Experience for more details.

Organize your resources with tags

Tags are now available for you to add to your resources, like Cloud Object Storage, to help you manage resources and find the resources that are the most relevant to you. For example, if you have hundreds of resources and you want to differentiate between ones that are paid the same way, you could tag them with `costcenter:location01`. Or, if you have a team that is working on a couple of resources repeatedly, you can use something like `team-blue`. You can also filter the My resources page by tags to quickly organize and find the resources that you need. For more information, see Working with tags and Platform Tagging on the Enhanced IBM Cloud Platform.

Get accurate monthly costs with the cost estimator

To help you decide and analyze what services you'd like to purchase, you can use the cost estimator. Now, you can go through the console and select each service you'd like to have, and add all of the costs in an easy to use tool. You can even enter projected data usages, lookups per second, writes per second, and queries per second to get a more accurate estimation of your monthly expenditures. You can use the cost estimator with each catalog service you select, or you can click the Cost Estimator icon in the console menu to get a summary of your estimated costs. For more information, see Estimating your costs.

01 November 2018

Updated global location names

As IBM Cloud continues to expand our global availability footprint, we're updating our location naming structure to better support an understandable, consistent hierarchy of geographies, regions, and data centers around the world. If you're familiar with our current global regions, you'll recognize names like US South and Sydney. We're aligning these location names to the names of the city in which the data centers physically exist.

For now, the programmatic IDs are not changing, so there's no impact from an API perspective. The following table shows the old and new location names. For more information and a comprehensive list of data centers and regions, see Service availability.

Previous Location Display Name	New Location Display Name	Code
US South	Dallas	us-south
US East	Washington DC	us-east
United Kingdom	London	eu-gb
Germany	Frankfurt	eu-de
Sydney	Sydney	au-syd
AP North	Tokyo	jp-tok

New location names

October 2018

30 October 2018

Assign account management access to others

With IBM Cloud Identity and Access Management (IAM), you can delegate common tasks that you complete as an account administrator to another user in your account. By creating an access policy on one or all of the available account management services, you can easily delegate responsibilities such as inviting and removing users, managing access groups, managing service IDs, maintaining private catalog services, and even monitoring billing and tracking usage. There are four individual account management services and an all services option that you can use to set up access policies:

- User Management for inviting and removing users
- IAM Access Groups for creating, editing, deleting, updating, and assigning access
- IAM Identity Service for viewing, creating, deleting, and assigning access to service IDs and associated API keys across the account
- Global resource catalog for viewing private catalog offerings and updating the metadata and visibility for the offerings
- All account management services for access to each of the individual account management service options based on the assigned role as well as access to billing and usage tracking.

For more information on the tasks that a user can do based on which account management service they have a policy on and which role they are assigned, see Example platform management roles and actions for account management services. For more information about this new feature, see the Introducing More Flexibility and Control for IBM Cloud Account Management Services Access blog post.

July 2018

17 July 2018

Searching for resources

You can search for resources from anywhere in the IBM Cloud console. Type the name of a resource in the search field in the console menu bar. Press the Forward Slash key (/) to activate the search.

12 July 2018

Dynamically add federated users to access groups

You can create dynamic rules to automatically add federated users to access groups based on specific identity attributes. When your users log in with a federated ID, the data from the identity provider dynamically maps your users to an access group based on the rules that you set. For more information, see Creating dynamic rules for access groups.

June 2018

01 June 2018

Protect your service IDs and API keys

To avoid a situation where your service ID or API key is deleted causing an outage or disruption, you can lock service IDs and API keys by using the UI or CLI. Locking a service ID also prevents any access policies from being changed, deleted, or assigned as well as any API keys associated with the service ID from being created or deleted. For more information, see Locking a service ID and Locking an API key.

May 2018

31 May 2018

Upgrade your Lite account to a Subscription account

You can now upgrade your Lite account to a Subscription account directly from the IBM Cloud console. With a Subscription account, you can use both platform and infrastructure offerings, and take advantage of discounted pricing by making a monthly spending and term commitment. You can also avoid surprises with fixed billing on a monthly payment schedule, but with the flexibility to order more or less based on your needs. For more information, see FAQS for billing and usage.

15 May 2018

IBM Cloud CLI rebranding

The IBM Cloud CLI commands changed from `bluemix` and `bx` to `ibmcloud`. However, you can still use the `bluemix` and `bx` CLI commands until they are removed later. There is no short name now, just the full name `ibmcloud`.

02 May 2018

Multi-factor authentication for your account

Multi-factor authentication (MFA) adds an extra layer of security to your account by requiring all users to provide a time-based one-time passcode in addition to their standard IBMid and password during login. This is also commonly known as two-factor authentication (2FA). MFA is enabled per account, and once it is turned on, all users in the account are required to log in by using the extra security measure. For more information, see the IBM Cloud Platform now adds support for Multi-Factor Authentication blog post.

April 2018

03 April 2018

Assign access quickly by using access groups

Do you want to be able to assign access quickly by using the least number of policies possible? Now you can with access groups. Group a set of users and service IDs together and assign a single policy that applies to all members of the group. By using access groups, you can limit the time that you spend managing access to the users and service IDs in your account.

December 2017

15 December 2017

Cloud Foundry Service US East region

A new US East data center is now available in Washington, DC. You can reach this new region by using the `us-east.cloud.ibm.com` endpoint. For details about the services that are available for purchase in this new region, see Services by region.

14 December 2017

Support for resources in the European Union

If your services and data centers are located in Europe, IBM Cloud now offers extra capabilities to protect your data in the European Union. You can request that support is provided by customer success teams that are located in Europe. This support is available 24 hours a day, 7 days a week. See

Enabling the EU supported option and Requesting support for resources in the European Union for more information.

November 2017

28 November 2017

Withdrawal of support for TLS 1.0 and 1.1

On 1 March 2018 IBM Cloud will withdraw support for TLS 1.0 and TLS 1.1 across many of our cloud products and services as part of our commitment to offering a cloud that is secure to the core and in alignment with industry best practices for security and data privacy.

16 November 2017

A new way to organize resources within your account

Resource groups are a new way for you to create customizable groupings of account resources, and access to the group and the resources within it are managed by using Identity and Access Management (IAM). Everyone starts out with a default resource group. You can rename this resource group and add new service instances to it as you create them from the catalog.

For users with a Pay-As-You-Go or Subscription account, you can create extra resource groups to make managing quota and viewing billing usage for a set of resources easier. You can also group resources to make it easier for you to assign users access to more than one service at a time. To learn more about working with resource groups for your account, see Managing resource groups.

Updates for IBM Cloud IAM

The introduction of resource groups within your IBM Cloud account provides a new way for you to assign access. Users and service IDs can be assigned access to all services within a resource group, enabling you to quickly assign access to more than one resource at a time. You can also customize access for each user or service ID by assigning access to just some services within a resource group, or you choose to assign access to individual resources down to the service instance level. For more information about the features that you can take advantage of by using IAM, see What features does IAM provide?

Customize your dashboard view

You can view and manage all the resources in your account from your dashboard in the IBM Cloud console. And now, you can set filters to customize your view. For example, you can filter by resource group to view the specific resources in a resource group. You can also filter by region or Cloud Foundry space.

02 November 2017

Support Center

We now have the new Support Center where you can search for information, post questions to our developer community, and manage tickets. Go to **Support > Support Center** in the IBM Cloud console menu bar.

October 2017

31 October 2017

Introducing IBM Cloud

Bluemix is now IBM Cloud. Besides rolling out our new name, nothing changes. You can still easily build and run your apps and services as always. Check out the IBM Cloud blog for more details.

Lite account

A Lite account is our new account type that gives you access to try select services for free with no time restrictions. This new account also includes usage tracking and efficiency features to help you better manage your resources. To learn more about what's available, see Account types.

06 October 2017

Identity and Access Management application authentication feature

Identity and Access Management (IAM) now supports service IDs, which you can think of as identities that can be used for apps to authenticate with your IBM Cloud services. Instead of using individual user credentials, a Service ID can be created with an associated API key and access permissions in the form of a service policy that is assigned to the Service ID in order for you to control the level of access for any application authenticating with that ID.

For more information about the benefits of this feature and how to get started, see the Introducing IBM Cloud IAM Service IDs and API Keys .

July 2017

27 July 2017

IBM Cloud global catalog

Expanding on the last console update to manage your public regions from a single location in the console, IBM Cloud now has a global catalog, making the process of selecting and deploying items that you select from the catalog a more streamlined process. Regardless of the region that you select in the console, you can now see all services that are available across all public regions from your catalog. Once you select a tile from the catalog, you can see which regions the service is available in, and select where you want to deploy it. For more information about the latest updates to the catalog, see A global IBM Cloud catalog makes building things easier .

May 2017

23 May 2017

IBM Cloud console updates

You can now manage your public regions from a single location through the updated IBM Cloud console. The region selector offers you streamlined access to your resources, and other enhancements include higher availability and improved performance.

01 May 2017

Identity and access management

With the latest updates and improvements, IBM Cloud account owners or administrators can now use a new unified access control UI to take advantage of the following capabilities:

- Manage users' fine-grained access to Kubernetes services and other services as they adopt the new access control features
- Assign service policies and Cloud Foundry roles to users within their organizations

Additionally, IBM Cloud platform users can create, delete, and list API keys associated with their user IDs. And platform users can use those API keys to authenticate when using APIs or CLIs.

Lastly, we enhanced our unified user management capability to ensure that in a linked IaaS-PaaS account, users are managed in a unified way with no need to add users separately in the SoftLayer Customer Portal or the IBM Cloudconsole.

April 2017

13 April 2017

Navigation design changes for IBM Cloud docs

> With this navigation update, we think you'll understand how content is better organized throughout our docs, and will be able to find relevant content more efficiently. With fewer nested layers of content, you won't have to dig around to find the documentation you need to be successful with IBM Cloud.

Try out IBM Cloud, for free

Looking to try out IBM Cloud®? Create an account and start building proof of concepts (POCs) with the many components available in IBM Cloud. You can try Lite and Free service plans to explore IBM Cloud at no cost while learning how to work in the cloud, use Watson, and more. This quick start guide is intended to help you get up and running on IBM Cloud without having to think about costs until you're ready.

Before you begin

Go to the IBM Cloud console and create an account. You're asked to enter your credit card information to secure your account and verify your identity. There are no costs that are associated with signing up, and you can try out IBM Cloud for free. You pay only for billable services that you choose to use, with no long-term contracts or commitments.

You're set up with a Pay-As-You-Go account, and you can access the full IBM Cloud catalog, including all Lite and Free plans. You receive a $200 credit to help get you started. You can use the $200 credit on IBM Cloud products that are used in the first 30 days.

Step 1: Explore the catalog

Explore the catalog for services that are free to use by filtering for Lite and Free pricing plans. You can work on your projects worry free, without the risk of generating a bill.

1. Go to the catalog.
2. Select the **Lite** and **Free** pricing plan options.

Step 2: Create an instance

Create an instance of product that includes a free Lite plan or a Free tier pricing plan.

1. Select the tile from the catalog after reviewing the filtered list of products.
2. Enter any required information to create the instance.
3. Click **Create**.

Step 3: Next steps: Check out the tutorials

Check out our tutorials for Lite plans for detailed steps about using IBM Cloud services that provide free Lite plans for you to implement common patterns based on best practices and proven technologies at no cost.

Navigating the IBM Cloud console

The IBM Cloud® console is the user interface that you use to manage all your IBM Cloud resources. You can create a free account, log in, access documentation, access the catalog, view pricing information, get support, or check the status of IBM Cloud components. After you log in, the menu bar contains a **Navigation Menu** icon and more links.

Using the console

When you log in to IBM Cloud, your dashboard is displayed, which shows widgets that summarize the status of your account. If you're interested in customizing your dashboard, see Working with scoped dashboards.

Use the following options to navigate to general areas of the console:

Browse available products

Use the **Catalog** link to explore over 350 products that offer options for compute, networking, security management, end-to-end developer solutions, and more. Use the tabs to filter the catalog to quickly access *deployable architectures*, IBM products, Cloud essentials, and more.

Find help when you need it

Click the **Help** icon > **Docs** to access the product documentation.

Get support when something's not working as expected

Click the **Help** icon > **Support center** to go to the Support Center page.

Manage account preferences and more

From the **Manage** menu, you can access your account, billing and usage, and Identity and Access Management options.

Quickly access a browser-based shell environment

Click the **IBM Cloud Shell** icon to open a browser-based shell environment that you can use to work with your IBM Cloud resources.

Estimate costs for your cloud deployments

Click the **Cost estimator** icon to open the cost estimator.

Stay up to date with notifications

Click the **Notifications** icon to view and control all incidents, maintenance, and announcements that are likely to affect your account.

Customize your profile and more

Click the **Avatar** icon to access your profile, guided tours, console theme options, and more.

In addition to the console, command-line interfaces (CLIs), APIs, and SDKs are available for interacting with you cloud account and resources. Terraform support is also available through use of the IBM Cloud Provider plug-in for managing cloud resources at enterprise scale through templates and scripting.

Can't find what you're looking for?

The IBM Cloud platform recently consolidated services and features to provide you with a more simplified and customized experience. Services and areas of the console are now unified into the following hubs: Infrastructure, Containers, Automation, Databases, Observability, and Security. The goal of these changes is to bundle together related services to make it easier to find, deploy, and use them. Use our mapping and guidance to see what's changed in the IBM Cloud catalog and console navigation to find what you're looking for.

Streamlined and dynamic IBM Cloud catalog

In addition to the new navigation changes, we updated the discoverability and management of our catalog. The new IBM Cloud catalog makes it easy to discover IBM products such as managed services, preconfigured software, professional services, and even partner services. You can navigate by using the new tabs in the catalog to explore what each category consists of.

- Get your job done more efficiently with expertly-designed solutions. Cloud essentials is your main source for exploring related products that are tailored to specific industries in a cohesive way.
- Use deployable architectures to understand how pre-built compositions of products work together to help solve common business problems. Deployable architectures accelerate your innovation and reduce risk across complex enterprise workloads. Industry solutions with configurations that previously took months to achieve are now available within hours.

Simplified access to your workloads

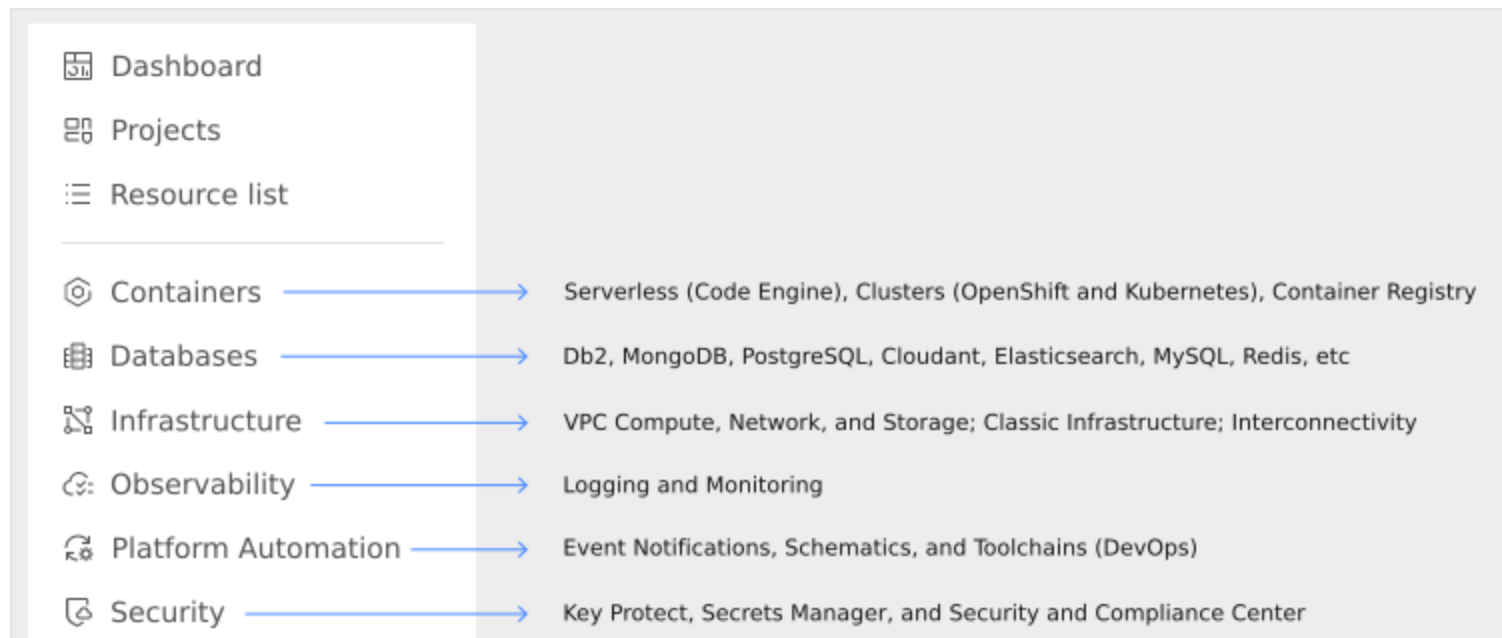

Updated navigation for the IBM Cloud console

From the **Navigation Menu** ☰, you can access areas of IBM Cloud that are focused on specific use cases and industries. Use the following options to explore the menu:

Containers

The previous **Kubernetes**, **OpenShift**, **Code Engine**, and **Container Registry** menu options are now nested within **Containers**.

Databases

All of your database needs are now centralized into one location. From **Databases**, you can create and manage relational databases (Databases for MySQL, Db2, Databases for PostgreSQL, and Databases for EDB); auxiliary databases (IBM Db2 Warehouse SaaS, Databases for etcd, and Databases for Redis); and non-relational databases (Databases for Elasticsearch, Databases for MongoDB, and IBM Cloudant).

Infrastructure

The previous **Classic Infrastructure**, **VPC Infrastructure**, and **Interconnectivity** menu options are now nested within **Infrastructure**.

Observability

The **Observability** option still provides access to **Logging**, **Monitoring**, and **Activity tracking**.

Platform Automation

The previous **DevOps** and **Schematics** options are now nested within **Platform Automation**, along with Event Notifications.

Security

Security and Compliance Center is now integrated with IBM Key Protect and Secrets Manager and can be accessed from the **Security** menu option.

Getting help from the AI assistant

IBM Cloud®'s AI assistant, which is powered by IBM®'s watsonx, is designed to help you learn about working in IBM Cloud and building solutions with the available catalog of offerings.

The AI assistant is designed as a retrieval-augmented generation (RAG) implementation that is running against IBM® large language models (LLMs) using IBM's watsonx and the IBM Cloud docs.

IBM Cloud does not gather any personal identifiable information when you are using the AI feature and it is not connected to or associated with your account. When you query the AI assistant, it is not being trained based on user input and does not associate any queries with your identity or account. IBM endeavors to develop AI in accordance with its published principles of trust and transparency. For more information, see AI Ethics.

Note: The AI assistant supports all national languages supported by IBM Cloud including Brazilian Portuguese, English, French, German, Japanese, Korean, Italian, Spanish, Simplified Chinese, and Traditional Chinese. To get responses from the AI assistant in one of the supported languages, ask your question in one of the languages and get an answer back in that same language.

Using the AI assistant in the IBM Cloud console

No matter where you are working in the IBM Cloud console, the AI assistant is available to help with your questions about IBM Cloud so that you can learn more about the platform or the services that you're using.

You can start the AI assistant from the IBM Cloud console header from the Help menu anytime that you want to ask a question, and it follows along with you in the console until you close it. With each question that you ask, reference citations from the IBM Cloud docs are provided.

1. From the IBM Cloud console menu bar, click the **Help** icon > **AI assistant**.
2. Ask questions about working in your account, working with products and offerings from the catalog, and more to help you stay productive in IBM Cloud.
3. Use the supported documentation references that are returned to have access to additional information and the ability to verify responses.

Important: Content that is generated in the AI assistant might include mistakes or be incorrect.

Submitting feedback for generated responses

You can provide feedback for each response from the AI assistant by using the **Good response** icon and **Bad response** icon options. You can choose to just mark the provided answer as a good response or bad response, and you can provide additional feedback by selecting from a common set of qualities that describe why you chose the rating and add additional feedback in the form of a comment. Your feedback is not used to train or enhance AI models.

Clearing your chat history

All of the questions that you ask and the generated responses are retained in the AI assistant until you decide that you want to clear the history. To clear the history, open the AI assistant, and then click **Open menu** icon > **Clear history**.

Closing the chat

The AI assistant follows along with you in the console until you close it. You can move it to position it wherever you'd like on the page. When you're ready to close it, click the **Close** icon . Your history is saved even when you close the AI assistant, so you can see your previous questions and answers the next time that you open it.

Using the AI assistant in the IBM Cloud CLI

You can also ask questions in the IBM Cloud CLI.

1. Log in with the `ibmcloud login` command. If you are logging in with a federated ID, use the `--sso` option to authenticate with a one-time passcode, or use the `--apikey` option to authenticate with an API key.
2. Ask a question with the `ibmcloud assist` command.

```
ibmcloud assist "How do I update the CLI?"
```

For more information about using the AI assistant in the IBM Cloud CLI, see General IBM Cloud CLI (ibmcloud) commands.

Limitations

The AI assistant has the following limitations:

- Questions must be limited to 500 characters or less.
- AI-generated responses are limited to information from the IBM Cloud docs and the LLM.
- Actions like creating resources or making changes within your account can't be completed by the AI assistant.
- Answers from the previous query and response are not used to answer subsequent queries.
- If you encounter any styling issues in the console, use the **Open menu** icon ☰ > **Clear history** option, and refresh the page.

Tips for creating effective prompts

When you ask the AI assistant a question, you are creating a prompt that is used to generate an answer. To ensure that you get the best response, review the following tips:

Always include context

: The more specific that you can be by including the right details, the better the response is. For example, if you're asking about a dashboard within the console for a specific service, you want to make sure to include the name of that service.

Ask questions that can be answered by the AI assistant

: The AI assistant is limited to answering questions about IBM Cloud. This means that your questions should remain about IBM Cloud and its offerings.

Be clear and concise

: You want to include enough detail to get your question answered correctly. But including too much detail or being too verbose in your prompt might result in an incorrect answer.

Be kind and polite

: Everyone wants to be treated kindly, even AI assistants.

Rephrase your question when needed

: If you don't get the answer that you expected, sometimes you need to rework your prompt to provide more or less detail or use different terms and synonyms to get the right response. For example, if you are looking for CLI instructions instead of console instructions, you might need to specify that you want information about a specific command to complete your task.

Here are some examples of questions that you might ask:

- What's the difference between Pay-As-You-Go and Subscription accounts?
- How can I increase my quotas for VPC resources?
- How should I organize my secrets?
- Help me learn more about Power Virtual Server

Privately connecting to IBM Cloud services

An increased focus on security is required by customers that use cloud-based services for production workloads. For many customers, accessing services in a secure manner is not only a sensible corporate policy, but in some cases required by compliance regulations. IBM® has enhanced the connectivity options for customers who require isolated connectivity options for their workloads by providing the options of virtual private endpoints (VPE) for Virtual Private Cloud (VPC) and IBM Cloud® service endpoints.

To interact with the product APIs, you can use public endpoints, VPE for VPC, or service endpoints.

- Public endpoints: You can connect to resources in all regions in your account over the IBM Cloud public network.
- VPE for VPC: After you create private endpoints in your VPC, you can connect by using a private IP address that's accessible only from your VPC. This option is available for VPC users.
- Service endpoints: After you enable virtual routing and forwarding (VRF) and turn on service endpoints for your account, you can connect by using a private IP address that's accessible only through the IBM Cloud private network. This option is available for classic infrastructure users.

By using VPE for VPC or service endpoints, you can privately connect to IBM Cloud platform services and service products in the catalog that support this type of private connectivity.

Tip: Go to the API docs for the service that you want to connect to and see the Endpoint URLs section to view the endpoints for that service.

Virtual private endpoints for VPC

VPE for VPC provides private connectivity to select IBM Cloud services originating from the VPC network of your choosing, without traversing the public backbone. All connectivity is contained within IBM Cloud. By using a service that supports VPE for VPC, you can connect to the service from your VPC network by assigning the IP address of your choosing, which is allocated from a subnet within your VPC.

VPEs are virtual IP interfaces that are bound to an endpoint gateway created on a per service or service instance basis, depending on the service operation model. The endpoint gateway is a virtualized function that scales horizontally, is redundant and highly available, and spans all availability zones of your VPC. Endpoint gateways enable communications from virtual server instances in your VPC to IBM Cloud services, with all traffic going over the IBM Cloud backbone. VPE for VPC gives you the experience of controlling all the private addressing within your cloud.

For more information, see About virtual private endpoint gateways.

IBM Cloud service endpoints

For classic infrastructure users, the cloud service endpoints feature provides the option of using private routes to connect to IBM Cloud services. These private routes are not accessible or reachable over the internet. You can protect your data from threats from the public network and logically extend your private network. This capability enables an enterprise with strict security requirements to have confidence in moving workloads to the IBM public cloud.

For more information, see Enabling virtual routing and forwarding (VRF) and service endpoints.

Encrypting data with your own keys

IBM Cloud® encrypts all data in transit and at rest. Additionally, you can configure IBM Cloud services to encrypt your data at rest with your own keys, and then monitor the events around the lifecycle of the encryption keys with IBM Cloud® Activity Tracker.

IBM Cloud supports multiple encryption options, whether you're looking for a solution that supports the bring your own key (BYOK) or keep your own key (KYOK) functionality. For an in-depth look at the options to secure your data depending on your organization's needs, see Data security.

Bring your own keys

Many IBM Cloud services support data encryption by using customer-managed keys, also known as bring your own key (BYOK). The most common use case for BYOK is using IBM® Key Protect to bring your encryption keys to the cloud. Key Protect is a multi-tenant service using FIPS 140-2 Level 3 HSM. For a list of services that can be integrated with Key Protect, see Integrating services.

Keep your own keys

IBM Cloud also provides IBM Cloud Hyper Protect Crypto Services, which is a dedicated key management service and hardware security module (HSM). Hyper Protect Crypto Services features keep your own key (KYOK) encryption capabilities backed by FIPS 140-2 Level 4 certification. With this option, it makes the IBM public cloud the industry's most secure and open public cloud for business. For a list of services that can be integrated with Hyper Protect Crypto Services, see Integrating IBM Cloud services with Hyper Protect Crypto Services.

Auditing the lifecycle of your keys

You can use Activity Tracker to monitor the activity of your keys. The Activity Tracker service provides the framework and functions to monitor API calls to services on IBM Cloud and produces the evidence to comply with corporate policies and market industry-specific regulations. Events that are tracked by Activity Tracker are either global or regional, and global events, such as provisioning a service, are available through the global domain instance that is located in Frankfurt. Events that are generated by an instance of Key Protect or Hyper Protect Crypto Services are automatically forwarded to the IBM Cloud instance that is available in the same location.

See Provisioning an instance to configure your monitoring instance. Whether you're using Key Protect or Hyper Protect Crypto Services, you can track events like creating a key, deleting a key, rotating a key, and more:

- IBM Cloud Activity Tracker events for Key Protect
- IBM Cloud Activity Tracker events for Hyper Protect Crypto Services

Getting certified in IBM Cloud

Are you looking to get certified in IBM Cloud®? Whether your job role is technical or nontechnical, the IBM Center for Cloud Training has a no-cost, role-based learning path to help guide you toward achieving certification. Start or continue building your cloud knowledge and skills today with training and certification from IBM.

IBM Cloud Advocate

Complete your IBM Cloud Advocate training and gain a foundational knowledge of cloud computing, from core concepts, definition, and history to emerging trends. Areas of learning include cloud service and deployment models as well as key elements of cloud architecture.

IBM Cloud Technical Advocate

Training for all technical roles starts with IBM Cloud Technical Advocate. Learn how to discuss and use IBM Cloud to solve business and technology transformation challenges. This learning path is recommended for all technical job roles and features practice exercises in a virtual lab environment.

IBM Cloud Professional Architect

As an IBM Certified Professional Architect technical professional, you have the skills and knowledge to define and design complex technical solutions for large-scale environments. Extensive course updates reflect the latest technologies and best practices for IBM Cloud architecture.

IBM Cloud Advanced Architect

Designed for professionals with three or more years of experience working as cloud architects, IBM Certified Advanced Architect extends and validates your capabilities. This interactive curriculum offers a deeper exploration of the latest IBM technologies and best practices.

IBM Cloud Professional Developer

Gain the skills and knowledge that is required to build cloud services and large-scale business applications as an IBM Certified Professional Developer technical professional. Learn cloud development and operations (DevOps) functions by using IBM Cloud services and APIs, including AI, cloud databases, and cloud security.

IBM Cloud Associate Site Reliability Engineer (SRE)

As an IBM Certified Associate Site Reliability Engineer technical professional, you are able to manage enterprise workloads in IBM Cloud environments. You also develop skills and knowledge around incident management, monitoring, troubleshooting, operations, deployments, and security and compliance.

IBM Cloud Professional Site Reliability Engineer

Advance your skills with IBM Certified Professional Site Reliability Engineer (SRE) professional-level training. Renew your understanding of operations, software engineering, and systems administration. And learn the monitoring and incident management tools that are needed to manage enterprise workloads in IBM Cloud environments.

IBM Cloud Security Engineer Specialty

Learn how to anticipate threats and maintain your organization’s security environment with an IBM Cloud Security Engineer Specialty certification. This course includes modules on secure infrastructure and hybrid cloud connections, Cloud compute, Kubernetes services, and VMware solutions in IBM Cloud.

IBM Cloud Satellite Specialty

As an IBM Cloud Satellite Specialty technical professional, you can deploy and run applications across on-premises, edge computing, and public cloud environments. Help your clients bring the flexibility and agility of public cloud services to their secure on-premises data centers and gain a competitive advantage.

IBM Cloud for Financial Services Specialty

An IBM Cloud for Financial Services Specialty certification teaches you the business drivers, differentiators, and components of IBM Cloud for Financial Services. Learn how to deliver operational criteria and compliance controls for confidential compute environments.

IBM Cloud DevSecOps Specialty

With an IBM Cloud DevSecOps Specialty certification, you are able to integrate security in a continuous delivery pipeline and understand how to deploy apps faster and more securely.

IBM Cloud IBM Cloud for VMware Solutions Specialty

The IBM Cloud for VMware Solutions Specialty prepares a cloud professional to design and implement IBM Cloud for VMware solutions. See IBM Cloud IBM Cloud for VMware Solutions Specialty.

IBM Cloud IBM Cloud for SAP Specialty

The IBM Cloud for SAP Specialty learning plan is an extension to professional-level, role-based IBM Cloud learning paths. It prepares cloud professionals to recommend IBM Cloud architectures meeting SAP application requirements and assist in migrating SAP workloads to the IBM Cloud. Check out IBM Cloud IBM Cloud for SAP Specialty.

Submitting feedback

We want to hear from you! You can submit feedback for the IBM Cloud team on the documentation or the console. You can choose from a few different methods to provide feedback.

Providing a content update with a pull request

You can submit suggested changes to any documentation page by using the **Edit topic** link on each page. By following that link and editing the page, you can submit updates in a pull request for the content team to review and merge.

Complete the following process to submit a pull request to the content team:

1. Click **Edit topic** on the page that you want to update.
2. Click the **Edit this file** icon.
3. Update the file in the WYSIWYG editor.
4. Scroll to the end of the page and enter a brief description about your change. The extended description is optional.
5. Click **Propose file change**.
6. Provide details about your changes, and click **Create pull request**.

Opening an issue for the documentation

You can open an issue for the content team to address from any documentation page. Within your issue, provide the details about the improvement that you'd like to see. For example, you might report that the steps on the page don't match the UI or more details about a specific subject are needed.

To open an issue from a documentation page, complete the following steps:

1. Click the **Open doc issue** link.
2. Enter a title and complete the template.
3. Click **Submit new issue**.

The content team works to address the issues that you open. You can monitor the issue by saving the link after you click **Submit new issue**. As the issue gets resolved, the content team provides status updates as a comment on the issue.

Submitting a comment by using the feedback form

When you're completing a task in the console, you can find a feedback form from the console menu bar to submit a rave review, a suggestion, or identify an error.

1. Click the **Help** icon ⓘ > **Send feedback** from the console menu bar.
2. Enter a description about what you want to share with the IBM Cloud team.
3. Rate your experience on a scale from one to five stars, with five stars representing a great experience.
4. Click **Submit**.

Submitting ideas for product improvements

To submit ideas for IBM Cloud and services improvements, go to IBM Cloud Ideas.

Product availability and locations

Service and infrastructure availability by location

IBM Cloud® makes it easier for you to implement, host, and scale services, infrastructure, and apps so you can focus on your application logic and application design. IBM Cloud's global network of locations provides three tiers of regions: *multizone regions single-campus multizone regions, and data centers. To achieve low application latency, deploy your apps in a region near your customers. For details about the available IBM Cloud regions and data centers for specific products, review the following sections.*

IBM provides SDKs and APIs for all services that are general availability. Check out the reference docs in the API & SDK reference library.

Not all services and infrastructure are available for purchase in every IBM Cloud location.

Services

Some services are available to purchase in a location, but that service's data might be hosted in a different location. The following table shows the services that are provided by IBM. For the full list of resources that are available, see the catalog in the IBM Cloud console.

Services that are hosted globally create resources that operate across multiple locations. For example, with IBM Cloud Object Storage, you can choose to deploy data in a single data center, or even a combination of locations by selecting the endpoint where your application sends REST API requests. For more information about the global network of locations, see Locations for resource deployment.

Service	Dallas (`us-south`)	Washington DC (`us-east`)	Toronto (`ca-tor`)	Sao Paulo (`br-sao`)
API Connect	✓	✓	✓	✓
Analytics Engine	✓			
App Configuration	✓	✓	✓	
App ID	✓	✓	✓	✓
Auto Scale for VPC	✓	✓	✓	✓
Bare Metal Servers for VPC	✓	✓	✓	✓
Block Storage Snapshots for VPC	✓	✓	✓	✓
Block Storage for VPC	✓	✓	✓	✓
Client VPN for VPC	✓	✓	✓	✓
Cloud Activity Tracker				
Cloud Logs	✓	✓	✓	✓
Cloud Monitoring	✓	✓	✓	✓
Cloud Object Storage	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Cloudant	✓	✓	✓	✓
Cluster Network	✓	✓	✓	✓
Code Engine	✓	✓	✓	✓
Compose for Elasticsearch	✓	✓		

Compose for MongoDB	✓	✓		
Compose for PostgreSQL	✓	✓		
Compose for RabbitMQ	✓	✓		
Compose for Redis	✓	✓		
Compose for etcd	✓	✓		
Container Registry	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Continuous Delivery	✓	✓	✓	✓
Cost and Asset Management	✓			
DNS Services	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Data Engine (previously SQL Query)	✓			
Data Product Hub	✓			
Data Replication	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Data Virtualization	✓			
DataStage	✓		✓	
Databases for EDB	✓	✓	✓	✓
Databases for Elasticsearch	✓	✓	✓	✓
Databases for MongoDB	✓	✓	✓	✓
Databases for MySQL	✓	✓	✓	✓
Databases for PostgreSQL	✓	✓	✓	✓
Databases for Redis	✓	✓	✓	✓
Databases for etcd	✓	✓	✓	✓
Db2	✓	✓	✓	✓
Db2 Warehouse	✓	✓		✓
Dedicated Host for VPC	✓	✓	✓	✓
Direct Link Connect	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Direct Link Dedicated	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Enterprise Application Service		✓		
Event Notifications	✓		✓	

Event Streams	✓	✓	✓	✓
File Storage for VPC	✓	✓	✓	✓
Floating IP for VPC	✓	✓	✓	✓
Flow Logs for VPC	✓	✓	✓	✓
Historical Instrument Analytics	✓			
Hyper Protect Crypto Services	✓	✓	✓	✓
IBM Cloud Backup for VPC	✓	✓	✓	✓
IBM Cloud Pak for Data	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
IBM Cloud Reservations for VPC	✓	✓	✓	✓
IBM Cognos Dashboard Embedded	✓			
IBM Knowledge Catalog	✓			
IBM Match 360 with Watson	✓			
IBM Verify				
IBM watsonx Code Assistant	✓			
Image Service for VPC	✓	✓	✓	✓
Instrument Analytics	✓			
Internet Services	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Internet of Things Platform				
Investment Portfolio	✓			
Key Protect	✓	✓	✓	✓
Knowledge Studio	✓	✓		
Kubernetes Service	✓	✓	✓	✓
Load Balancer for VPC	✓	✓	✓	✓
Log Analysis				
MQ	✓	✓		
Messages for RabbitMQ	✓	✓	✓	✓
Multi Volume Snapshots for VPC	✓	✓	✓	✓
Natural Language Understanding	✓	✓		

Netezza Performance Server	✓	✓	✓	
Network ACL	✓	✓	✓	✓
OpenPages	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Partner with Technology Expert Labs	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Placement Groups for VPC	✓	✓	✓	✓
Planning Analytics	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Portfolio Optimization	✓			
Power Virtual Server Network	✓	✓	✓	✓
Power Virtual Server Shared Processor Pool	✓	✓	✓	✓
Power Virtual Server Snapshot	✓	✓	✓	✓
Power Virtual Server Virtual Machine	✓	✓	✓	✓
Power Virtual Server Volume	✓	✓	✓	✓
Predictive Market Scenarios	✓			
Private Path Service for VPC	✓	✓	✓	✓
Public Gateway	✓	✓	✓	✓
Qiskit Runtime		✓		
Real-Time Payments	✓			
Red Hat OpenShift on IBM Cloud	✓	✓	✓	✓
SSH Key for VPC	✓	✓	✓	✓
Satellite	✓	✓	✓	✓
Satellite Infrastructure Service				
Schematics	✓	✓		
Secrets Manager	✓	✓	✓	✓
Secure Gateway	✓	✓		
Security Group for VPC	✓	✓	✓	✓
Security and Compliance Center	✓		✓	
Security and Compliance Center Workload Protection	✓	✓	✓	✓
Simulated Historical Instrument Analytics	✓			

Simulated Instrument Analytics	✓			
Software Instance	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Speech to Text	✓	✓		
Streaming Analytics	✓	✓		
Subnet	✓	✓	✓	✓
Text to Speech	✓	✓		
Toolchain	✓	✓	✓	✓
Transit Gateway	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
VCF as a Service - Cloud Director Site	✓	✓	✓	
VCF as a Service - Virtual Data Center	✓	✓	✓	
VMware Solutions	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
VPN for VPC	✓	✓	✓	✓
Virtual Network Interface	✓	✓	✓	✓
Virtual Private Cloud	✓	✓	✓	✓
Virtual Private Endpoint for VPC	✓	✓	✓	✓
Virtual Server for VPC	✓	✓	✓	✓
Watson Discovery	✓	✓		
Workspace for Power Virtual Server	✓			
watsonx Assistant	✓	✓		
watsonx Orchestrate	✓	✓	✓	
watsonx.ai Runtime	✓			
watsonx.ai Studio	✓			
watsonx.data	✓	✓		
watsonx.governance	✓			

Service availability - Americas

Service	London (`eu-gb`)	Frankfurt (`eu-de`)	Madrid (`eu-es`)
API Connect	✓	✓	
Analytics Engine		✓	

App Configuration	✓	✓	
App ID	✓	✓	
Auto Scale for VPC	✓	✓	✓
Bare Metal Servers for VPC	✓	✓	✓
Block Storage Snapshots for VPC	✓	✓	✓
Block Storage for VPC	✓	✓	✓
Client VPN for VPC	✓	✓	✓
Cloud Activity Tracker			
Cloud Logs	✓	✓	✓
Cloud Monitoring	✓	✓	✓
Cloud Object Storage	Hosted Globally	Hosted Globally	Hosted Globally
Cloudant	✓	✓	✓
Cluster Network	✓	✓	✓
Code Engine	✓	✓	✓
Compose for Elasticsearch	✓	✓	
Compose for MongoDB	✓	✓	
Compose for PostgreSQL	✓	✓	
Compose for RabbitMQ	✓	✓	
Compose for Redis	✓	✓	
Compose for etcd	✓	✓	
Container Registry	Hosted Globally	Hosted Globally	Hosted Globally
Continuous Delivery	✓	✓	✓
Cost and Asset Management			
DNS Services	Hosted Globally	Hosted Globally	Hosted Globally
Data Engine (previously SQL Query)		✓	
Data Product Hub			
Data Replication	Hosted Globally	Hosted Globally	Hosted Globally
Data Virtualization	✓	✓	

DataStage		✓	
Databases for EDB	✓	✓	✓
Databases for Elasticsearch	✓	✓	✓
Databases for MongoDB	✓	✓	✓
Databases for MySQL	✓	✓	✓
Databases for PostgreSQL	✓	✓	✓
Databases for Redis	✓	✓	✓
Databases for etcd	✓	✓	✓
Db2	✓	✓	
Db2 Warehouse	✓	✓	
Dedicated Host for VPC	✓	✓	✓
Direct Link Connect	Hosted Globally	Hosted Globally	Hosted Globally
Direct Link Dedicated	Hosted Globally	Hosted Globally	Hosted Globally
Enterprise Application Service			
Event Notifications	✓	✓	✓
Event Streams	✓	✓	✓
File Storage for VPC	✓	✓	✓
Floating IP for VPC	✓	✓	✓
Flow Logs for VPC	✓	✓	✓
Historical Instrument Analytics			
Hyper Protect Crypto Services	✓	✓	✓
IBM Cloud Backup for VPC	✓	✓	✓
IBM Cloud Pak for Data	Hosted Globally	Hosted Globally	Hosted Globally
IBM Cloud Reservations for VPC	✓	✓	✓
IBM Cognos Dashboard Embedded	✓		
IBM Knowledge Catalog	✓	✓	
IBM Match 360 with Watson			
IBM Verify		✓	

IBM watsonx Code Assistant			
Image Service for VPC	✓	✓	✓
Instrument Analytics			
Internet Services	Hosted Globally	Hosted Globally	Hosted Globally
Internet of Things Platform		✓	
Investment Portfolio			
Key Protect	✓	✓	✓
Knowledge Studio	✓	✓	
Kubernetes Service	✓	✓	✓
Load Balancer for VPC	✓	✓	✓
Log Analysis			
MQ	✓	✓	
Messages for RabbitMQ	✓	✓	✓
Multi Volume Snapshots for VPC	✓	✓	✓
Natural Language Understanding	✓	✓	
Netezza Performance Server	✓	✓	
Network ACL	✓	✓	✓
OpenPages	Hosted Globally	Hosted Globally	Hosted Globally
Partner with Technology Expert Labs	Hosted Globally	Hosted Globally	Hosted Globally
Placement Groups for VPC	✓	✓	✓
Planning Analytics	Hosted Globally	Hosted Globally	Hosted Globally
Portfolio Optimization			
Power Virtual Server Network	✓	✓	✓
Power Virtual Server Shared Processor Pool	✓	✓	✓
Power Virtual Server Snapshot	✓	✓	✓
Power Virtual Server Virtual Machine	✓	✓	✓
Power Virtual Server Volume	✓	✓	✓
Predictive Market Scenarios			

Private Path Service for VPC	✓	✓	✓
Public Gateway	✓	✓	✓
Qiskit Runtime		✓	
Real-Time Payments			
Red Hat OpenShift on IBM Cloud	✓	✓	
SSH Key for VPC	✓	✓	✓
Satellite	✓	✓	✓
Satellite Infrastructure Service		✓	
Schematics	✓	✓	
Secrets Manager	✓	✓	✓
Secure Gateway	✓	✓	
Security Group for VPC	✓	✓	✓
Security and Compliance Center		✓	✓
Security and Compliance Center Workload Protection	✓	✓	✓
Simulated Historical Instrument Analytics			
Simulated Instrument Analytics			
Software Instance	Hosted Globally	Hosted Globally	Hosted Globally
Speech to Text	✓	✓	
Streaming Analytics	✓	✓	
Subnet	✓	✓	✓
Text to Speech	✓	✓	
Toolchain	✓	✓	✓
Transit Gateway	Hosted Globally	Hosted Globally	Hosted Globally
VCF as a Service - Cloud Director Site	✓	✓	
VCF as a Service - Virtual Data Center	✓	✓	
VMware Solutions	Hosted Globally	Hosted Globally	Hosted Globally
VPN for VPC	✓	✓	✓
Virtual Network Interface	✓	✓	✓

Virtual Private Cloud	✓	✓	✓	
Virtual Private Endpoint for VPC	✓	✓	✓	
Virtual Server for VPC	✓	✓	✓	
Watson Discovery	✓	✓		
Workspace for Power Virtual Server				
watsonx Assistant	✓	✓		
watsonx Orchestrate	✓			
watsonx.ai Runtime	✓	✓		
watsonx.ai Studio	✓	✓		
watsonx.data	✓	✓		
watsonx.governance		✓		

Service availability - Europe

Service	Sydney (au-syd)	Tokyo (jp-tok)	Osaka (jp-osa)	Chennai (in-che)
API Connect	✓	✓		
Analytics Engine				
App Configuration	✓	✓	✓	
App ID	✓	✓	✓	
Auto Scale for VPC	✓	✓	✓	
Bare Metal Servers for VPC	✓	✓	✓	
Block Storage Snapshots for VPC	✓	✓	✓	
Block Storage for VPC	✓	✓	✓	
Client VPN for VPC	✓	✓	✓	
Cloud Activity Tracker				
Cloud Logs	✓	✓	✓	
Cloud Monitoring	✓	✓	✓	
Cloud Object Storage	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Cloudant	✓	✓	✓	✓
Cluster Network	✓	✓	✓	

Code Engine	✓	✓	✓	
Compose for Elasticsearch	✓			
Compose for MongoDB	✓			
Compose for PostgreSQL	✓			
Compose for RabbitMQ	✓			
Compose for Redis	✓			
Compose for etcd	✓			
Container Registry	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Continuous Delivery	✓	✓	✓	
Cost and Asset Management				
DNS Services	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Data Engine (previously SQL Query)				
Data Product Hub	✓			
Data Replication	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Data Virtualization		✓		
DataStage	✓			
Databases for EDB	✓	✓	✓	✓
Databases for Elasticsearch	✓	✓	✓	✓
Databases for MongoDB	✓	✓	✓	✓
Databases for MySQL	✓	✓	✓	✓
Databases for PostgreSQL	✓	✓	✓	✓
Databases for Redis	✓	✓	✓	✓
Databases for etcd	✓	✓	✓	✓
Db2	✓	✓		
Db2 Warehouse	✓	✓		
Dedicated Host for VPC	✓	✓	✓	
Direct Link Connect	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Direct Link Dedicated	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally

Enterprise Application Service				
Event Notifications	✓	✓	✓	
Event Streams	✓	✓	✓	✓
File Storage for VPC	✓	✓	✓	
Floating IP for VPC	✓	✓	✓	
Flow Logs for VPC	✓	✓	✓	
Historical Instrument Analytics				
Hyper Protect Crypto Services		✓		
IBM Cloud Backup for VPC	✓	✓	✓	
IBM Cloud Pak for Data	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
IBM Cloud Reservations for VPC	✓	✓	✓	
IBM Cognos Dashboard Embedded				
IBM Knowledge Catalog	✓	✓		
IBM Match 360 with Watson				
IBM Verify				
IBM watsonx Code Assistant				
Image Service for VPC	✓	✓	✓	
Instrument Analytics				
Internet Services	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Internet of Things Platform				
Investment Portfolio				
Key Protect	✓	✓	✓	
Knowledge Studio	✓	✓		
Kubernetes Service	✓	✓	✓	✓
Load Balancer for VPC	✓	✓	✓	
Log Analysis				
MQ				
Messages for RabbitMQ	✓	✓	✓	✓

Multi Volume Snapshots for VPC	✓	✓	✓	
Natural Language Understanding	✓	✓		
Netezza Performance Server		✓		
Network ACL	✓	✓	✓	
OpenPages	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Partner with Technology Expert Labs	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Placement Groups for VPC	✓	✓	✓	
Planning Analytics	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Portfolio Optimization				
Power Virtual Server Network	✓	✓	✓	
Power Virtual Server Shared Processor Pool	✓	✓	✓	
Power Virtual Server Snapshot	✓	✓	✓	
Power Virtual Server Virtual Machine	✓	✓	✓	
Power Virtual Server Volume	✓	✓	✓	
Predictive Market Scenarios				
Private Path Service for VPC	✓	✓	✓	
Public Gateway	✓	✓	✓	
Qiskit Runtime				
Real-Time Payments				
Red Hat OpenShift on IBM Cloud	✓	✓		✓
SSH Key for VPC	✓	✓	✓	
Satellite	✓	✓	✓	
Satellite Infrastructure Service				
Schematics				
Secrets Manager	✓	✓	✓	
Secure Gateway	✓			
Security Group for VPC	✓	✓	✓	
Security and Compliance Center				

Security and Compliance Center Workload Protection	✓	✓	✓	
Simulated Historical Instrument Analytics				
Simulated Instrument Analytics				
Software Instance	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
Speech to Text	✓	✓		
Streaming Analytics				
Subnet	✓	✓	✓	
Text to Speech	✓	✓		
Toolchain	✓	✓	✓	
Transit Gateway	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
VCF as a Service - Cloud Director Site	✓	✓		
VCF as a Service - Virtual Data Center	✓	✓		
VMware Solutions	Hosted Globally	Hosted Globally	Hosted Globally	Hosted Globally
VPN for VPC	✓	✓	✓	
Virtual Network Interface	✓	✓	✓	
Virtual Private Cloud	✓	✓	✓	
Virtual Private Endpoint for VPC	✓	✓	✓	
Virtual Server for VPC	✓	✓	✓	
Watson Discovery	✓	✓		
Workspace for Power Virtual Server				
watsonx Assistant	✓	✓		
watsonx Orchestrate				
watsonx.ai Runtime	✓	✓		
watsonx.ai Studio	✓	✓		
watsonx.data	✓	✓		
watsonx.governance	✓			

Service availability - Asia Pacific

Classic infrastructure

Classic infrastructure services are available to be deployed in data centers.

Americas

The following classic infrastructure resources are available in North and South America. Match the resource row with the column representing the data center. A check mark indicates the resource is available in that location.

Service	Montreal 01	Toronto 01
Bare Metal Server for Classic	✓	✓
Citrix NetScaler VPX	✓	✓
Cloud HSM	✓	✓
Cloud Load Balancer	✓	✓
Content Delivery Network	✓	✓
Direct Link Connect on Classic		
Direct Link Dedicated Hosting on Classic		✓
Direct Link Dedicated on Classic	✓	✓
Direct Link Exchange on Classic		
File Storage for Classic	✓	✓
FortiGate Security Appliance 10Gbps		✓
Gateway Appliance	✓	✓
Hardware Firewall	✓	✓
IBM Cloud Backup for Classic	✓	✓
IPSec VPN	✓	✓
Juniper vSRX	✓	✓
SSL Certificates	✓	✓
Secondary Subnets	✓	✓
Virtual Router Appliance	✓	✓

Americas infrastructure availability - Canada

Service	Dallas 09	Dallas 10	Dallas 12	Dallas 13	Dallas 14	San Jose 03	San Jose 04	Washington DC 04	Washington DC 06	Washington DC 07
Bare Metal Server for Classic	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Citrix NetScaler VPX	✓	✓	✓	✓		✓	✓	✓	✓	✓
Cloud HSM	✓	✓	✓	✓			✓	✓	✓	✓

Cloud Load Balancer	✓	✓	✓	✓		✓	✓	✓	✓	✓
Content Delivery Network	✓	✓	✓	✓		✓	✓	✓	✓	✓
Direct Link Connect on Classic										
Direct Link Dedicated Hosting on Classic	✓	✓	✓	✓		✓		✓		
Direct Link Dedicated on Classic	✓	✓	✓	✓		✓	✓	✓	✓	✓
Direct Link Exchange on Classic										
File Storage for Classic	✓	✓	✓	✓		✓	✓	✓	✓	✓
FortiGate Security Appliance 10Gbps	✓	✓	✓	✓		✓	✓	✓	✓	✓
Gateway Appliance	✓	✓	✓	✓		✓	✓	✓	✓	✓
Hardware Firewall	✓	✓	✓	✓		✓	✓	✓	✓	✓
IBM Cloud Backup for Classic	✓	✓	✓	✓		✓	✓	✓	✓	✓
IPSec VPN	✓	✓	✓	✓		✓	✓	✓	✓	✓
Juniper vSRX	✓	✓	✓	✓		✓	✓	✓	✓	✓
SSL Certificates	✓	✓	✓	✓		✓	✓	✓	✓	✓
Secondary Subnets	✓	✓	✓	✓		✓	✓	✓	✓	✓
Virtual Router Appliance	✓	✓	✓	✓		✓	✓	✓	✓	✓

Americas infrastructure availability - United States

Service	Sao Paulo 01
Bare Metal Server for Classic	✓
Citrix NetScaler VPX	✓
Cloud HSM	
Cloud Load Balancer	✓
Content Delivery Network	✓
Direct Link Connect on Classic	
Direct Link Dedicated Hosting on Classic	✓

Direct Link Dedicated on Classic	✓
Direct Link Exchange on Classic	
File Storage for Classic	✓
FortiGate Security Appliance 10Gbps	
Gateway Appliance	✓
Hardware Firewall	✓
IBM Cloud Backup for Classic	✓
IPSec VPN	✓
Juniper vSRX	✓
SSL Certificates	✓
Secondary Subnets	✓
Virtual Router Appliance	✓

Americas infrastructure availability - Brazil

Europe

The following classic infrastructure resources are available in Europe. Match the resource row with the column representing the data center. A check mark indicates the resource is available in that location.

Service	Amsterdam 03
Bare Metal Server for Classic	✓
Citrix NetScaler VPX	✓
Cloud HSM	✓
Cloud Load Balancer	✓
Content Delivery Network	✓
Direct Link Connect on Classic	
Direct Link Dedicated Hosting on Classic	
Direct Link Dedicated on Classic	✓
Direct Link Exchange on Classic	
File Storage for Classic	✓
FortiGate Security Appliance 10Gbps	✓
Gateway Appliance	✓

Hardware Firewall	✓
IBM Cloud Backup for Classic	✓
IPSec VPN	✓
Juniper vSRX	✓
SSL Certificates	✓
Secondary Subnets	✓
Virtual Router Appliance	✓

Europe infrastructure availability - Netherlands

Service	Frankfurt 02	Frankfurt 04	Frankfurt 05
Bare Metal Server for Classic	✓	✓	✓
Citrix NetScaler VPX	✓	✓	✓
Cloud HSM	✓		
Cloud Load Balancer	✓	✓	✓
Content Delivery Network	✓	✓	✓
Direct Link Connect on Classic			
Direct Link Dedicated Hosting on Classic	✓		
Direct Link Dedicated on Classic	✓	✓	✓
Direct Link Exchange on Classic			
File Storage for Classic	✓	✓	✓
FortiGate Security Appliance 10Gbps	✓	✓	✓
Gateway Appliance	✓	✓	✓
Hardware Firewall	✓	✓	✓
IBM Cloud Backup for Classic	✓	✓	✓
IPSec VPN	✓	✓	✓
Juniper vSRX	✓	✓	✓
SSL Certificates	✓	✓	✓
Secondary Subnets	✓	✓	✓
Virtual Router Appliance	✓	✓	✓

Europe infrastructure availability - Germany

Service	London 02	London 04	London 05	London 06
Bare Metal Server for Classic		✓	✓	✓
Citrix NetScaler VPX	✓	✓	✓	✓
Cloud HSM	✓	✓		
Cloud Load Balancer	✓	✓	✓	✓
Content Delivery Network	✓	✓	✓	✓
Direct Link Connect on Classic				
Direct Link Dedicated Hosting on Classic	✓			✓
Direct Link Dedicated on Classic	✓	✓	✓	✓
Direct Link Exchange on Classic				
File Storage for Classic	✓	✓	✓	✓
FortiGate Security Appliance 10Gbps	✓	✓		
Gateway Appliance	✓	✓	✓	✓
Hardware Firewall	✓	✓	✓	✓
IBM Cloud Backup for Classic	✓	✓	✓	✓
IPSec VPN	✓	✓	✓	✓
Juniper vSRX	✓	✓	✓	✓
SSL Certificates	✓	✓	✓	✓
Secondary Subnets	✓	✓	✓	✓
Virtual Router Appliance	✓	✓	✓	✓

Europe infrastructure availability - United Kingdom

Service	Milan 01
Bare Metal Server for Classic	✓
Citrix NetScaler VPX	✓
Cloud HSM	
Cloud Load Balancer	✓
Content Delivery Network	✓
Direct Link Connect on Classic	
Direct Link Dedicated Hosting on Classic	

Direct Link Dedicated on Classic	✓
Direct Link Exchange on Classic	
File Storage for Classic	✓
FortiGate Security Appliance 10Gbps	
Gateway Appliance	✓
Hardware Firewall	✓
IBM Cloud Backup for Classic	✓
IPSec VPN	✓
Juniper vSRX	✓
SSL Certificates	✓
Secondary Subnets	✓
Virtual Router Appliance	✓

Europe infrastructure availability - Italy

Service	Paris 01
Bare Metal Server for Classic	✓
Citrix NetScaler VPX	✓
Cloud HSM	✓
Cloud Load Balancer	✓
Content Delivery Network	✓
Direct Link Connect on Classic	
Direct Link Dedicated Hosting on Classic	✓
Direct Link Dedicated on Classic	✓
Direct Link Exchange on Classic	
File Storage for Classic	✓
FortiGate Security Appliance 10Gbps	✓
Gateway Appliance	✓
Hardware Firewall	✓
IBM Cloud Backup for Classic	✓
IPSec VPN	✓

Juniper vSRX			✓
SSL Certificates			✓
Secondary Subnets			✓
Virtual Router Appliance			✓

Europe infrastructure availability - France

Service	Madrid 02	Madrid 04	Madrid 05
Bare Metal Server for Classic	✓	✓	✓
Citrix NetScaler VPX			
Cloud HSM			
Cloud Load Balancer			
Content Delivery Network	✓		
Direct Link Connect on Classic			
Direct Link Dedicated Hosting on Classic			
Direct Link Dedicated on Classic			
Direct Link Exchange on Classic			
File Storage for Classic			
FortiGate Security Appliance 10Gbps			
Gateway Appliance	✓	✓	✓
Hardware Firewall			
IBM Cloud Backup for Classic			
IPSec VPN			
Juniper vSRX			
SSL Certificates			
Secondary Subnets			
Virtual Router Appliance			

Europe infrastructure availability - Spain

Asia Pacific

The following classic infrastructure resources are available in Asia and the Pacific region. Match the resource row with the column representing the data center. A check mark indicates the resource is available in that location.

Service	Sydney 01	Sydney 04	Sydney 05

Bare Metal Server for Classic	✓	✓	✓
Citrix NetScaler VPX	✓	✓	✓
Cloud HSM	✓	✓	
Cloud Load Balancer	✓	✓	✓
Content Delivery Network	✓	✓	✓
Direct Link Connect on Classic			
Direct Link Dedicated Hosting on Classic	✓		✓
Direct Link Dedicated on Classic	✓	✓	✓
Direct Link Exchange on Classic			
File Storage for Classic	✓	✓	✓
FortiGate Security Appliance 10Gbps	✓		
Gateway Appliance	✓	✓	✓
Hardware Firewall	✓	✓	✓
IBM Cloud Backup for Classic	✓	✓	✓
IPSec VPN	✓	✓	✓
Juniper vSRX	✓	✓	✓
SSL Certificates	✓	✓	✓
Secondary Subnets	✓	✓	✓
Virtual Router Appliance	✓	✓	✓

Asia Pacific infrastructure availability - Australia

Service	Chennai 01
Bare Metal Server for Classic	✓
Citrix NetScaler VPX	✓
Cloud HSM	✓
Cloud Load Balancer	✓
Content Delivery Network	✓
Direct Link Connect on Classic	✓
Direct Link Dedicated Hosting on Classic	
Direct Link Dedicated on Classic	✓

Direct Link Exchange on Classic	
File Storage for Classic	✓
FortiGate Security Appliance 10Gbps	
Gateway Appliance	✓
Hardware Firewall	✓
IBM Cloud Backup for Classic	✓
IPSec VPN	✓
Juniper vSRX	✓
SSL Certificates	✓
Secondary Subnets	✓
Virtual Router Appliance	✓

Asia Pacific infrastructure availability - India

Service	Tokyo 02	Tokyo 04	Tokyo 05
Bare Metal Server for Classic	✓		✓
Citrix NetScaler VPX	✓	✓	✓
Cloud HSM	✓	✓	
Cloud Load Balancer	✓	✓	✓
Content Delivery Network	✓	✓	✓
Direct Link Connect on Classic		✓	
Direct Link Dedicated Hosting on Classic	✓		
Direct Link Dedicated on Classic	✓	✓	✓
Direct Link Exchange on Classic			
File Storage for Classic	✓	✓	✓
FortiGate Security Appliance 10Gbps	✓		
Gateway Appliance	✓	✓	✓
Hardware Firewall	✓	✓	✓
IBM Cloud Backup for Classic	✓	✓	✓
IPSec VPN	✓	✓	✓
Juniper vSRX	✓	✓	✓

SSL Certificates	✓	✓	✓
Secondary Subnets	✓	✓	✓
Virtual Router Appliance	✓	✓	✓

Asia Pacific infrastructure availability - Japan

Service	Singapore 01
Bare Metal Server for Classic	
Citrix NetScaler VPX	✓
Cloud HSM	
Cloud Load Balancer	✓
Content Delivery Network	✓
Direct Link Connect on Classic	
Direct Link Dedicated Hosting on Classic	✓
Direct Link Dedicated on Classic	✓
Direct Link Exchange on Classic	
File Storage for Classic	✓
FortiGate Security Appliance 10Gbps	✓
Gateway Appliance	✓
Hardware Firewall	✓
IBM Cloud Backup for Classic	✓
IPSec VPN	✓
Juniper vSRX	✓
SSL Certificates	✓
Secondary Subnets	✓
Virtual Router Appliance	✓

Asia Pacific infrastructure availability - Singapore

Select services and enhancements coming soon

Information regarding potential future services is intended to outline our general direction, and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future services isn't a commitment, promise, or legal obligation to deliver any material, code, or functionality. Information about potential future services might not be incorporated into any contract.

The following list includes select services and enhancements that are coming soon:

- IBM Cloud Logs: Frankfurt, Madrid

- L4 and L40S GPU based virtual server profiles: Sao Paolo
- watsonx.governance: Frankfurt
- watsonx Code Assistant for Z: Frankfurt
- watsonx Code Assistant for Red Hat Ansible Lightspeed: Frankfurt
- IBM Cloud Logs: Dallas, London, Tokyo, Toronto, Sydney, Sao Paolo, Washington DC
- H100 GPU based virtual server profiles: London, Sydney, Tokyo
- Databases for MongoDB sharding with scale to 12 terabytes: Dallas, Frankfurt, London, Madrid, Tokyo, Toronto, Sydney, Sao Paolo, Washington DC
- Dedicated host for VPC - Reserved Instance: Dallas, Frankfurt, London, Madrid, Tokyo, Toronto, Sydney, Sao Paolo, Washington DC
- watsonx Code Assistant for Z: London, Tokyo
- watsonx Code Assistant for Red Hat Ansible Lightspeed: London, Tokyo

IBM Cloud region and data center locations for resource deployment

IBM Cloud® has a resilient global network of locations to host your highly available cloud workload. Resources in different locations are consolidated into an account-based billing and usage view. You can also deploy your workloads to the location that is nearest to your customers to achieve low latency connectivity. IBM Cloud provides *multizone regions (MZR)*, *single-campus multizone regions (SC-MZR)*, and classic *data centers* for classic infrastructure resources.

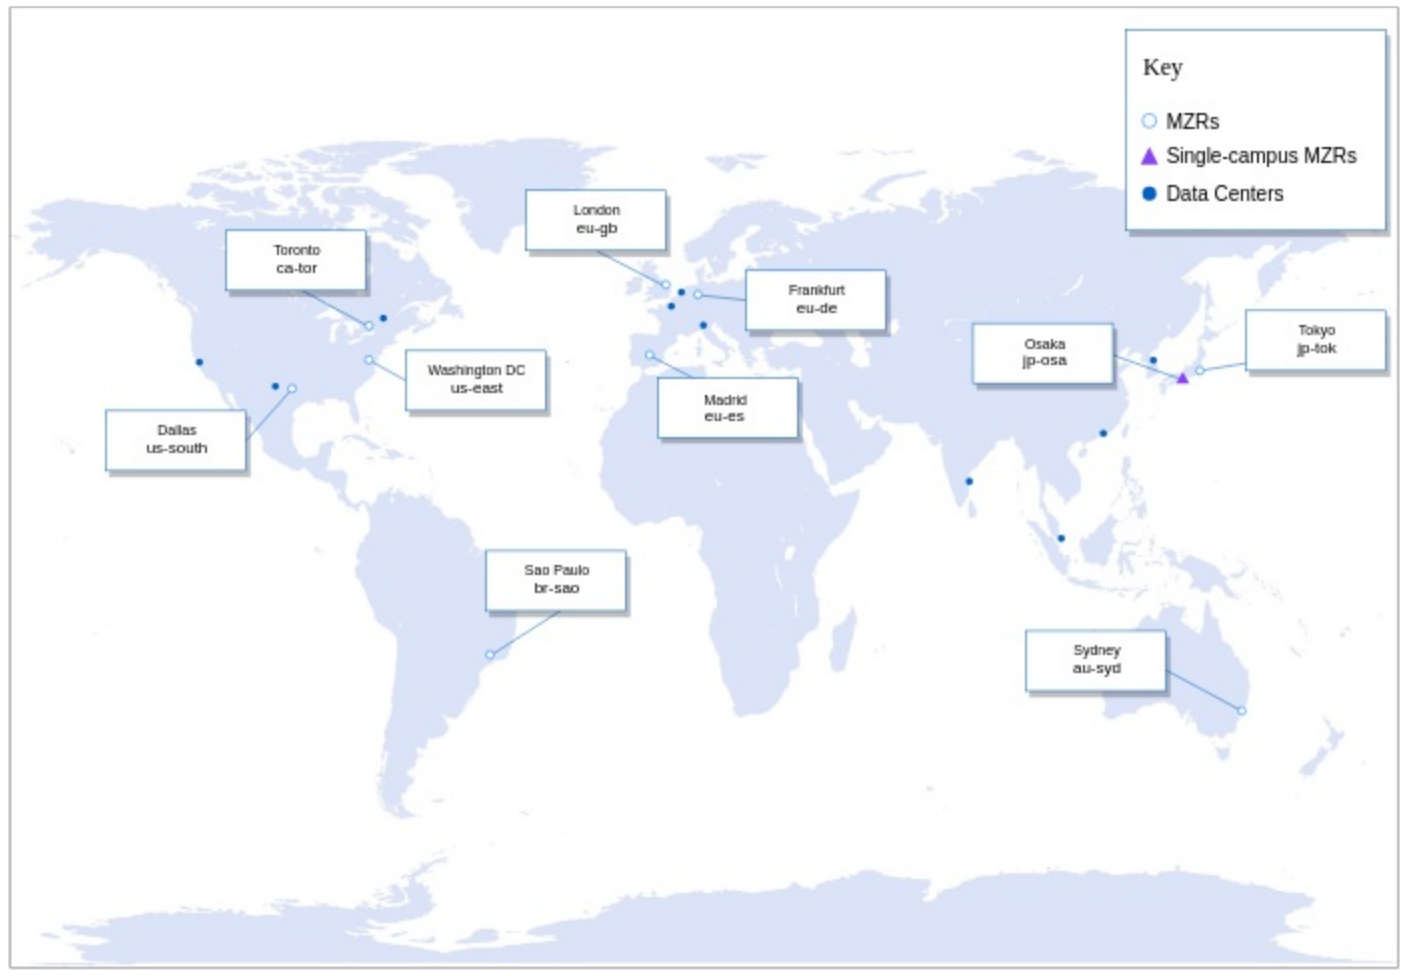

MZR and data center locations map

Note: This image is an artistic representation and does not reflect actual political or geographic boundaries.

Regions

IBM® offers two types of regions: MZRs and single-campus MZRs and both are considered an MZR. The underlying infrastructure in both types provides the same SLA. A region is an independent geographic territory that consists of one or more zones and is typically referred to by the metropolitan (metro) city area name like Dallas or London.

Each *zone* within the region assists with improved fault tolerance and decreased latency. A zone is identified by using two separate names. There is a zone name, for example `us-south-1` that is a logical identifier for a zone in the context of the current account. There is also a universal zone name that is the identifier for a zone that is consistent across IBM Cloud, for example `us-south-dal10-a`. The universal zone name provides the location specification for VPC resources by mapping the zone name to a physical location, such as a data center. Alternatively, the location for classic resources is not specified by zone and instead uses the specific data center within the region, such as `DAL10`. For more information about zone information specific to your account, see Zone mapping per account.

By distributing your workloads across three zones and consuming highly available regional cloud resources through virtual private endpoints, you can increase regional availability. Distributing a workload across multiple regions can provide higher availability and serve as the foundation for a disaster recovery plan. There are zonal, regional, and global cloud services that provide a consistent set of resources across regions. The IBM regional services are distributed across zones in an MZR and generally provide 99.99% (tier 3) availability.

Multizone regions

MZRs are composed of three or more data centers in multiple zones with independent power, cooling, and network connectivity to help ensure that failures in these components will be isolated to a single zone. MZRs provide low latency (< 2-milliseconds latency) and high bandwidth (> 1000 Gbps) connectivity within a zone.

Offering the highest level of redundancy and availability by leveraging three separate sites within a region, MZRs have a minimum distance of at least 1 mile between zones and exact distances vary by region. Zone-to-zone latency can be found in the network latency dashboards.

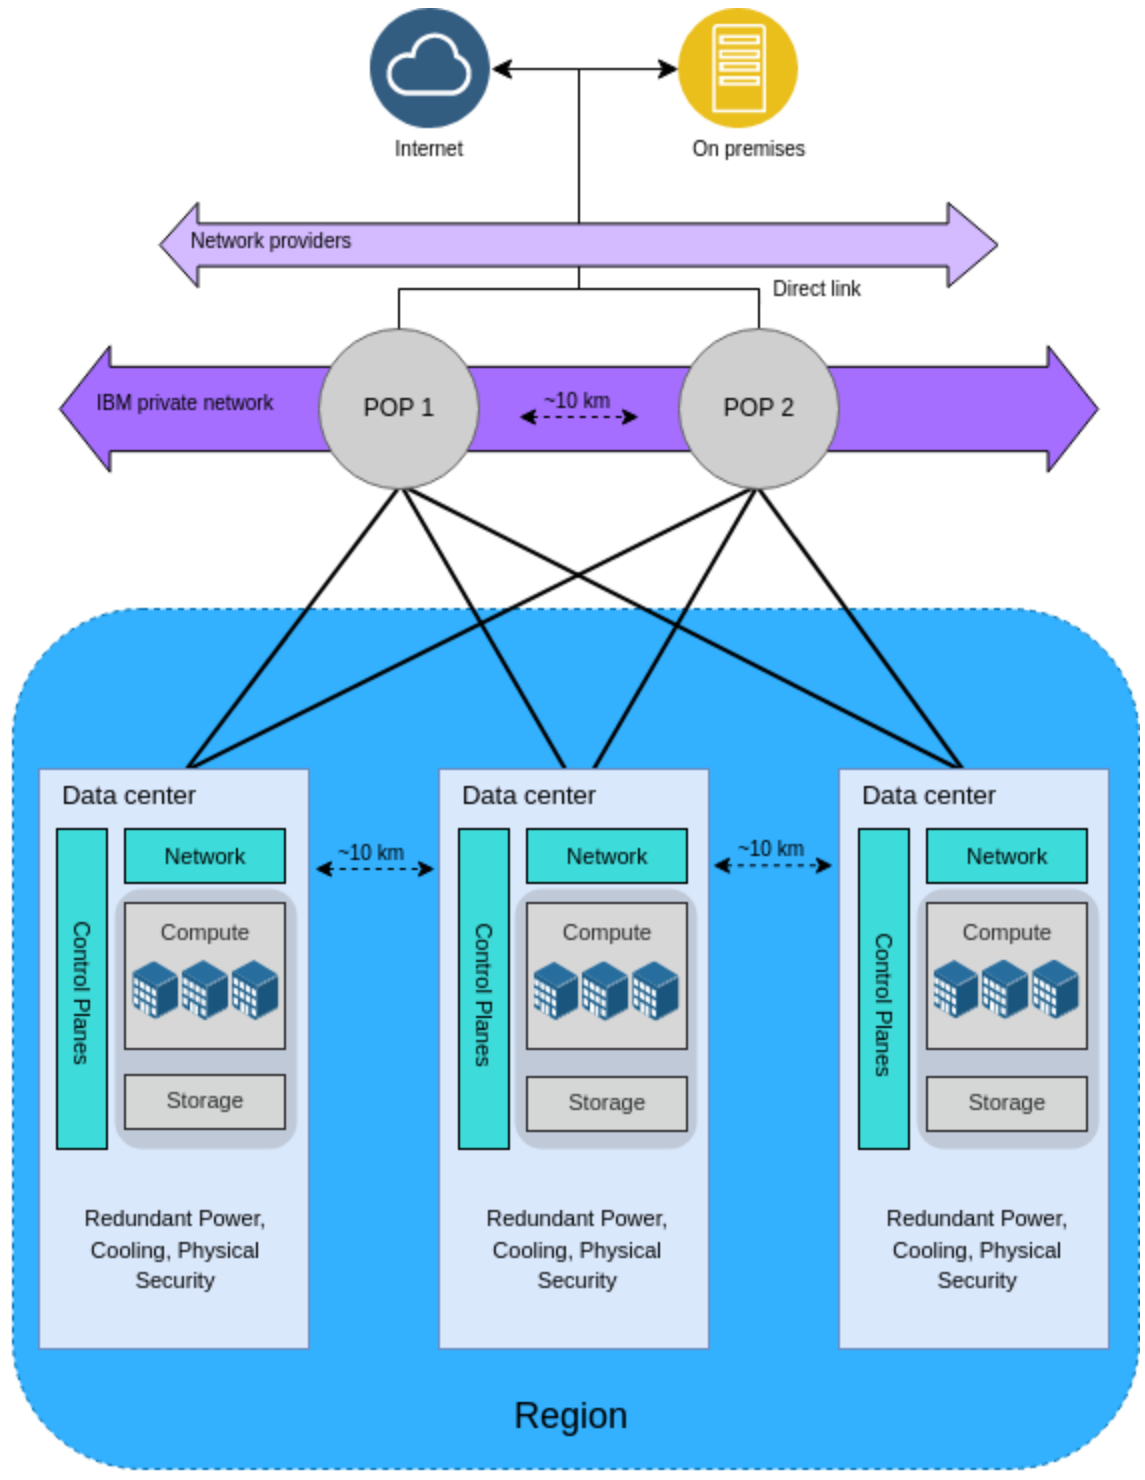

Multizone region (MZR)

MZRs support different types of compute for both VPC and classic infrastructure resources. The location of classic resources is specified by a data center while VPC resource locations are specified by the zone. For more information about the physical locations available for your account per region for VPC resources, see Zone mapping per account.

The following table lists the IBM Cloud MZR locations and zones for each.

Region	Zone
Dallas (`us-south`)	us-south-1 us-south-2 us-south-3
Sao Paulo (`br-sao`)	br-sao-1 br-sao-2 br-sao-3
Toronto (`ca-tor`)	ca-tor-1 ca-tor-2 ca-tor-3

Washington DC (`us-east`)	us-east-1 us-east-2 us-east-3

MZRs in North and South America

Region	Zone
Frankfurt (`eu-de`)	eu-de-1 eu-de-2 eu-de-3
London (`eu-gb`)	eu-gb-1 eu-gb-2 eu-gb-3
Madrid (`eu-es`)	eu-es-1 eu-es-2 eu-es-3

MZRs in Europe

Region	Zone
Sydney (`au-syd`)	au-syd-1 au-syd-2 au-syd-3
Tokyo (`jp-tok`)	jp-tok-1 jp-tok-2 jp-tok-3

MZRs in Asia Pacific

Note: If you're referencing a region when using the CLI, API, SDK, or Terraform, ensure that you're using the programmatic region name. For example, use `us-south` to target the Dallas (`us-south`) region.

Single-campus MZRs

Single-campus MZRs (SC-MZR) contain three zones in different sections of the same building or within multiple buildings on a campus where the power, cooling, networking, and physical security dependencies overlap but are not identical between any two zones. This setup ensures a level of continuous availability and survivability of any one system outage, planned or unplanned.

SLAs are maintained because the infrastructure is set up in a concurrently maintainable fashion so that a single failure does not affect all three zones in the same campus. This setup is ideal for services that support local users as it reduces latency or to support disaster recovery workloads.

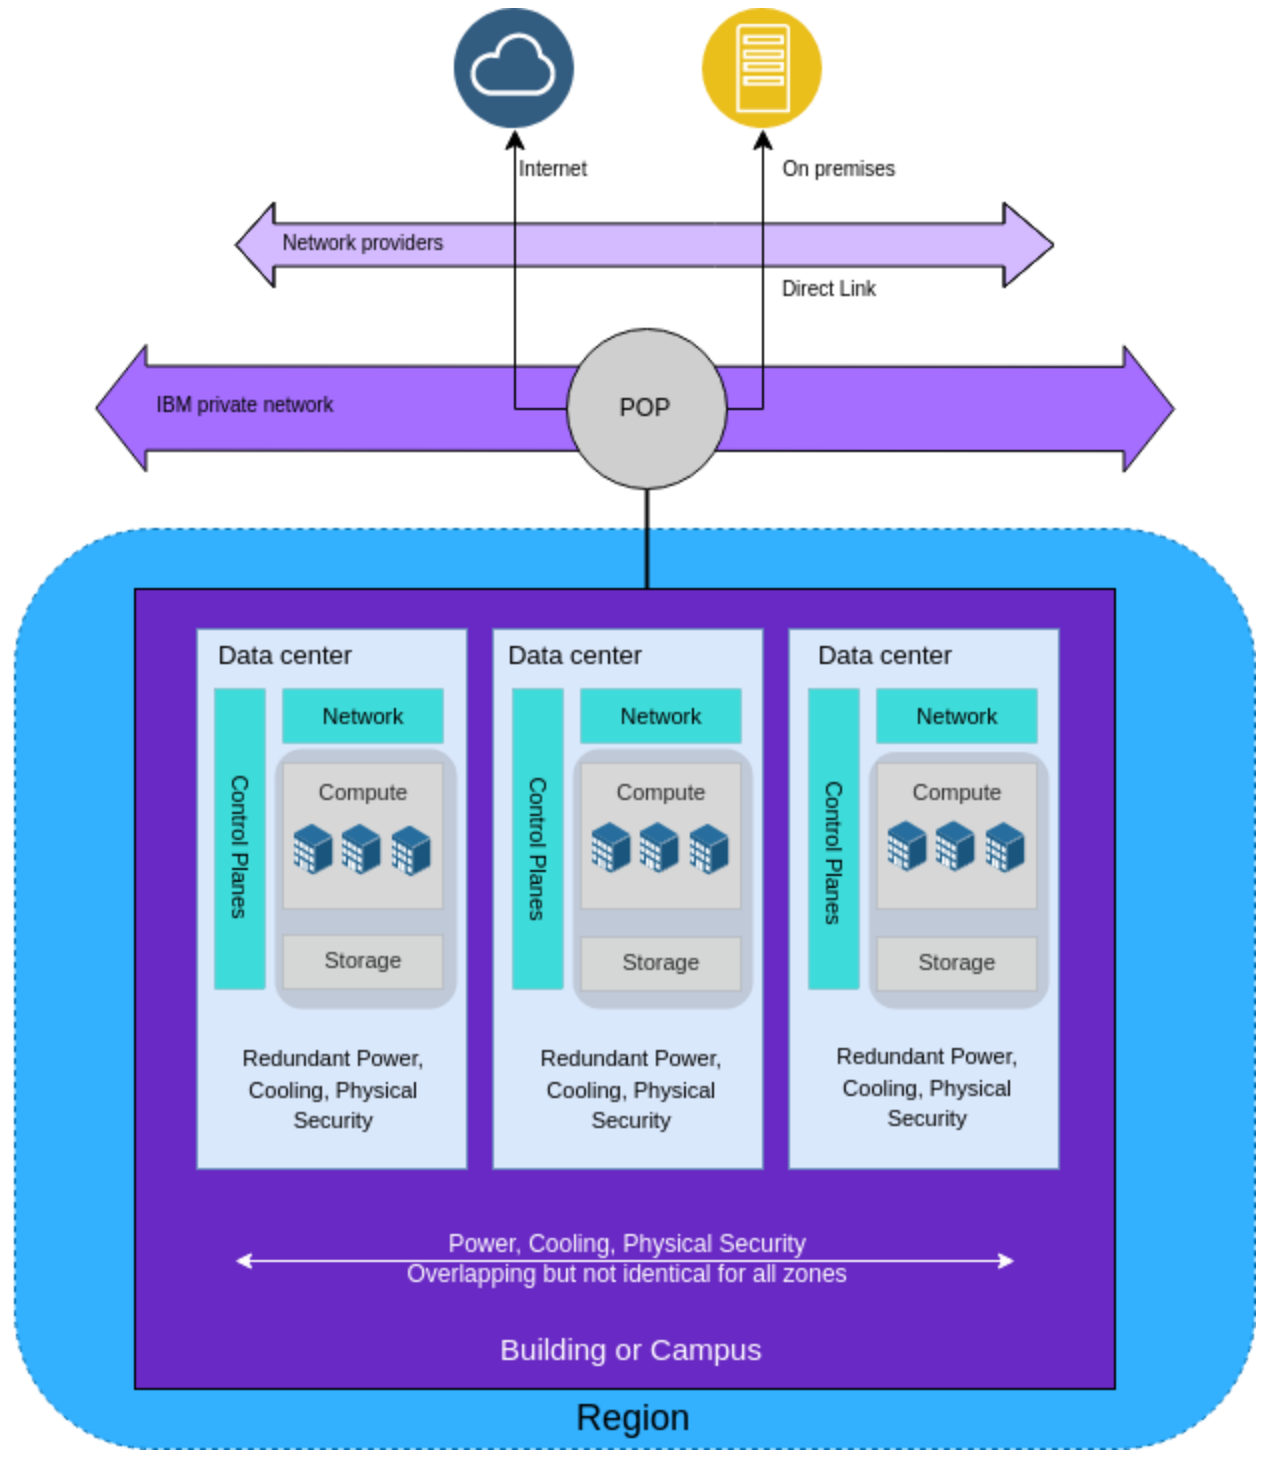

Single-campus MZR

The following table lists the SC-MZR locations that are available in IBM Cloud and the associated regions and zones.

Region	Zone
Osaka (`jp-osa`)	jp-osa-1 jp-osa-2 jp-osa-3

single-campus MZRs

Zone mapping per account

Within each region, there are three or more zones that are identified in the API, SDK, CLI, and Terraform by using a `regionname-number` syntax, for example `us-south-1`. Each IBM Cloud account has a zone mapping for each region that determines the relationship between the zone and the physical location. The zones map to a physical location, which is referred to by a universal zone name by using a `regionname-datacenter-letter` syntax, for example `us-south-dal10-a`.

The account zone mapping is established when the first VPC resource is created in the region, and it can't be changed. You can review the assigned zone mapping for an account on the VPC Infrastructure Overview page in the Endpoint section. You can also use the VPC API to list the mapping for your account.

Tip: Understanding your account's zone mapping is helpful if you're creating a mixed VPC and Power application, for example. You can create your VPC resources first, and then review your zone mapping to determine which universal zone the VPC resources are in so that you can ensure that the classic resources are created in the same physical location. Classic infrastructure and IBM® Power® Virtual Server services locations are specified by data center while the physical location for VPC resources are specified by the universal zone name.

The following table shows the available physical locations by using their universal zone name, associated data centers, and available *Point of Presence (PoP)* locations per MZR.

Region	Universal zone name	Data center	PoP

Dallas (`us-south`)	us-south-dal10-a us-south-dal12-a us-south-dal13-a us-south-dal14-a	DAL10 DAL12 DAL13 DAL14	DAL03 DAL04
Sao Paulo (`br-sao`)	br-sao-sao01-a br-sao-sao04-a br-sao-sao05-a	SAO01 SAO04 SAO05	SAO02 SAO03
Toronto (`ca-tor`)	ca-tor-tor01-a ca-tor-tor04-a ca-tor-tor05-a	TOR01 TOR04 TOR05	TOR02 TOR03
Washington DC (`us-east`)	us-east-wdc04-a us-east-wdc06-a us-east-wdc07-a	WDC04 WDC06 WDC07	WDC02 WDC05

MZR universal zone names - North and South America

Region	Universal zone name	Data center	PoP
Frankfurt (`eu-de`)	eu-de-fra02-a eu-de-fra04-a eu-de-fra05-a	FRA02 FRA04 FRA05	FRA01 FRA03
London (`eu-gb`)	eu-gb-lon04-a eu-gb-lon05-a eu-gb-lon06-a	LON04 LON05 LON06	LON01 LON03
Madrid (`eu-es`)	eu-es-mad02-a eu-es-mad04-a eu-es-mad05-a	MAD02 MAD04 MAD05	MAD01 MAD03

MZR universal zone names - Europe

Region	Universal zone name	Data center	PoP
Sydney (`au-syd`)	au-syd-syd01-a au-syd-syd04-a au-syd-syd05-a	SYD01 SYD04 SYD05	MEL02 PER01 SYD02 SYD03
Tokyo (`jp-tok`)	jp-tok-tok02-a jp-tok-tok04-a jp-tok-tok05-a	TOK02 TOK04 TOK05	TOK01 TOK03

MZR universal zone names - Asia Pacific

Note: If you're referencing a region when using the CLI, API, SDK, or Terraform, ensure that you're using the programmatic region name. For example, use `us-south` to target the Dallas (`us-south`) region.

The following table shows the available physical locations using their universal zone name, associated data centers, and available PoP locations per SC-MZR.

Region	Universal zone name	Data center	PoP
Osaka (`jp-osa`)	jp-osa-osa21-a jp-osa-osa22-a jp-osa-osa23-a	OSA21 OSA22 OSA23	OSA01

single-campus MZR universal zone name

Viewing resources by location

You can view all resources and locations from the Resource list page in the console. If you want to view and work with resources in a specific location,

expand the **Location** filter, and select a location from the list. By expanding a specific location, you can select to filter by regions, zones, or individual data centers.

Depending on the type of resource, you might be interested in only specific types of location data. For example, if you created a service or VPC infrastructure service, you can filter the Resource list page by the region and zone codes. However, if you're working with classic infrastructure or Power Virtual Server resources, the specific data center codes are the pertinent information for you.

For example, if you have resources that are deployed in the London 2 (eu-gb-2) zone, you can set filters to display only those resources in your resource list. Expand the **London** metro option, and the **London (eu-gb)** region option. Within that region, you can select from the list of available zones, such as **London 2 (eu-gb-2)**.

If you have a classic infrastructure resource that is deployed in a specific data center, you can identify the data center by the specific metro location and alphanumeric code. For example, use **Dallas** for the metro location and then **Dallas 10 (dal10)** for the data center.

You can also view resources that are deployed in Satellite locations, which are managed by an IBM Cloud metro or region and determines where the master of your Satellite control plane runs. For example, you might have a Satellite location that's managed by the Dallas metro. Expand the **Dallas** metro option, which includes your Satellite location, like `my-satellite-dal`. For more information about the metros and regions that manage Satellite locations, see Regions.

You might also want to display your resources that are located globally. The **Global** option means that only one logical, globally accessible instance of the service, independent of any region or zone, is published to customer workloads. These types of resources are accessible from a global endpoint.

As illustrated in the following graphic, a data center is a physical building that represents a zone that is located within a multizone region (MZR). An MZR is organized by its metro location. For example, London can encompass more than one grouping of data centers within an MZR. The graphic shows three zones in one MZR that work together in the instance that one of the data centers becomes unavailable. Zones are connected directly to each or through low latency links.

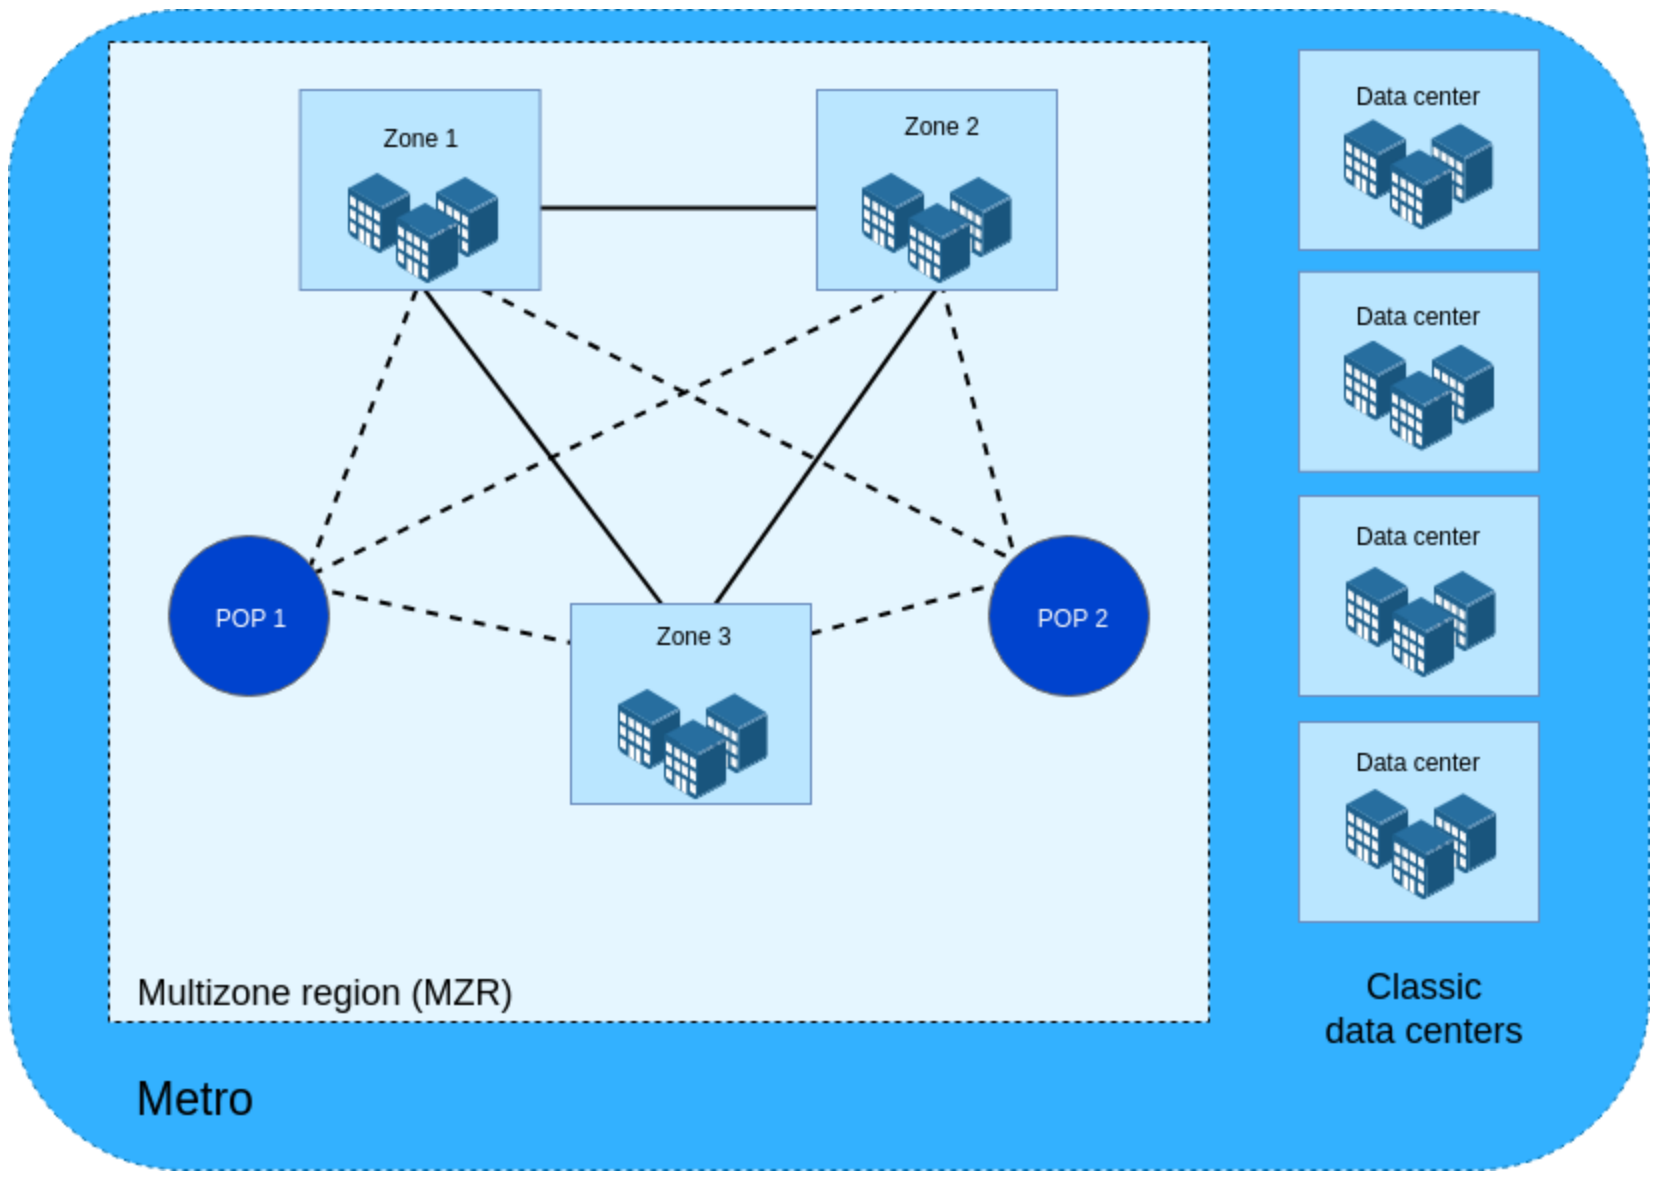

Location hierarchy

Classic data centers

In addition to selecting a region for your resource, you can select from a list of the IBM Cloud data centers, if you're working with classic infrastructure or Power Virtual Server resources.

Data centers host the power, cooling, compute, network, and storage resources used for services and apps. They don't provide isolation from multizones in a location.

Data centers are based on a POD architecture where each data center can have more than one POD, depending on the on-demand build out. Each POD consists of racks, servers, networks, and storage, along with backup power generators. Placing workload servers across PODs improves the availability.

See the following table for the specific code for each data center.

Data center	Code
Dallas 08 [1]	DAL08
Dallas 09	DAL09
Dallas 10	DAL10
Dallas 12	DAL12
Dallas 13	DAL13
Dallas 14	DAL14
Montreal 01	MON01
San Jose 03	SJC03
San Jose 04	SJC04
Sao Paulo 01	SAO01
Sao Paulo 04	SAO04
Sao Paulo 05	SAO05
Toronto 01	TOR01
Toronto 04	TOR04
Toronto 05	TOR05
Washington DC 03 [2]	WDC03
Washington DC 04	WDC04
Washington DC 06	WDC06
Washington DC 07	WDC07

Data centers in North and South America

Data center	Code
Amsterdam 03	AMS03
Frankfurt 02	FRA02
Frankfurt 04	FRA04
Frankfurt 05	FRA05
London 02	LON02
London 04	LON04
London 05	LON05

London 06	LON06
Madrid 02	MAD02
Madrid 04	MAD04
Madrid 05	MAD05
Milan 01	MIL01
Paris 01	PAR01

Data centers in Europe

Data center	Code
Chennai 01	CHE01
Osaka 21	OSA21
Osaka 22	OSA22
Osaka 23	OSA23
Singapore 01	SNG01
Sydney 01	SYD01
Sydney 04	SYD04
Sydney 05	SYD05
Tokyo 02	TOK02
Tokyo 04	TOK04
Tokyo 05	TOK05

Data centers in Asia Pacific

Note: The table includes certain data centers that are set to close soon. For the list of data centers that are closing, see Data center closures.

1. IBM Cloud for Government Learn more ↩
2. IBM Cloud for Government Learn more ↩

Service rollout policy

IBM Cloud® has a resilient global network of locations to host your highly available cloud workload. To ensure that the cloud infrastructure and services are consistent and stable across our deployment locations, we created best practices for our service catalog management. These best practices help us to accomplish rollouts in the most efficient manner and minimize business impact, costs, and risks. The following information describes our guidelines on when to expect or how to request that a service is available in your region.

This policy covers all IBM Cloud public MZRs, public single-campus MZRs and public data centers.

IBM® classifies our services deployed to our public locations as core or market-driven.

Core services

All IBM® multi-zone regions contain the following core services, which are the most basic and vital services that are needed for the majority of customer

workloads.

- IBM Cloud platform (console, CLI, Identity and Access Management, and global catalog)
- IBM Cloud® Virtual Private Cloud
    - IBM® Cloud Block Storage for Virtual Private Cloud
    - IBM Cloud® Virtual Servers for Virtual Private Cloud
    - Virtual Private Network (VPN) for VPC
    - IBM Cloud® Transit Gateway
    - Network Load Balancer for VPC
    - Application Load Balancer for VPC
    - Virtual Private Endpoint (VPE) for VPC
    - IBM Cloud® DNS Services
- IBM Cloud Object Storage
- IBM Cloud Databases for PostgreSQL
- IBM Key Protect for IBM Cloud
- IBM Cloud® Continuous Delivery
- IBM Cloud Container Registry
- IBM Cloud Kubernetes Service
- Red Hat OpenShift on IBM Cloud

Note: The IBM Cloud platform, including the console, CLI, Identity and Access Management, and global catalog, is a globally accessible instance that is independent of any region or zone. Global resources like the platform are accessible from a global endpoint.

Deployment Tiers

IBM identifies the following deployment tiers that can contain core services, market driven services or both.

Deployment tier	Core service	Market-driven services
MZR	✓	✓
Single campus MZR	✓	✓
Data center		✓

Deployment tiers and service types

Core service deployments

Adding new core service to existing MZRs

After a new core service is deployed in the first MZR and added to this IBM Cloud service rollout policy, the new core service will be deployed to all other MZRs within a period of 90 days.

Updating existing core services in existing MZRs

After a generally available update to an existing core service is deployed in the first MZR and documented in a release note, the same update will be deployed to all other MZRs within a period of 30 days.

Important: Not all hardware dependent profiles and features are available in all MZRs. If the service you want depends on such a profile or feature, contact IBM Cloud Sales for details on availability.

Note: Some services could deploy sooner.

Market-driven deployments

Market-driven services are deployed based on sufficient customer demand. To request that one of these services is available in your region, contact IBM Cloud Sales.

Deployments into any location other than an MZR are always market-driven.

Note: Market-driven classification covers any case other than those specified under the MZRs' description.

Dependency and service availability standards

For each service, you can review our dependency and availability standards.

- IBM Cloud services are available in several regions worldwide.
- IBM services can be deployed by using various models, and the SLA standards that apply to each model are outlined in the Service Level Agreements
- Follow our best practices to ensure your environment that integrates the use of the IBM Cloud services is as resilient as possible. Check out the resiliency documentation that is published in the architecture center.
- To build a solution that meets your needs, it is important for you to see the Shared responsibility matrix for more information on the scope of the services.
- For quick reference, we include links to the high availability documents for the services within our SLO description.

Dependency levels

Low-level implementation and details of services change based on your setup and configuration requirements. As such, we've provided a high-level explanation for the dependencies.

These levels might be considered layers or tiers in an architectural block diagram. Each layer might depend only on the layers below it, although there are some dependencies that might be fulfilled in the same layer in the case of some complex services.

Core services form the foundation for our cloud infrastructure. These services are required for a new MZR location to go live and are required to be updated across the entire footprint when there is a fundamental change.

Market-driven services are based on market demand and depend on the core services and some other services within this market-driven category. The other services are services that are needed to support regulation that applies in different regions or industries.

Important: Market-driven classification covers all services other than those specified under the core service description.

Service-to-service dependencies

The IBM Cloud backend services support all of the other services within our deployment units.

IBM Cloud services have a set of dependencies that are self-contained within the MZR deployment units. The rest of the service dependencies use global services that are served from cross regional locations. Any solution that is not hosted on an MZR might require an associated MZR that hosts the control plane services.

Note: If you are a current IBM Cloud customer and have a non-disclosure agreement with IBM, you can request the service dependency reports by going to Compliance support under PaaS compliance report and submit a request.

For the key dimensions of compute, storage, networking, and authentication or authorization, we indicate the following dependencies.

Control planes

Common dependencies for control planes are:

- The identity or platform data plane for authentication and authorization
- The audit tracking service
- Internal services that provide, for example, workflow, metadata storage, monitoring and logging
- Load balancers, VPN, and other network infrastructure services

Some control planes obviously have service-specific dependencies. For example, the compute control plane, when starting a bare metal or VM instance, depends on:

- Object Storage to retrieve the required system images
- Block Volumes control plane for provisioning and attaching the volumes

- Networking control plane to provision and attach NICs, subnets, and so on

Data planes

Core service data planes follow the general principle that each data plane is designed to have minimal dependencies in order to achieve high availability, fast time to diagnosis, and fast time to recovery. Therefore, the systems continue to function (with some limitations, for example cannot provision, reboot, or decom) even when there is a control plane outage.

Network

The networking data and control planes exist in every deployment unit.

Some services depend on others in the core, for example Compute services like bare metal and VM instances depend the Block Volumes data plane and the Networking data plane, and therefore will be impacted by the data planes for these services going down. But as stated before the control planes might suffer an outage and the Compute instances continue to function, if they do not require the associated service control plane (that is, increasing volume storage size, and so on)

Internal services

All client facing internal or external services depend on the identity and platform data plane for authentication and authorization. Control planes for monitoring, logging, and IBM Cloud CLI and data plane services for security services depend on the identity platform data plane. (Specific configurations might vary.)

Storage

The Object Storage data plane does not depend on Block Volumes or File Storage. Services that support backup and restore depend on Object Storage to operate. Therefore, those services require that the Block Volumes data and control planes exist in every deployment unit.

Related documents

Document and Link	Description
SLA	IBM Cloud SLA
Shared Responsibility Matrix	IBM's customer shared responsibility matrix
Availability of services	Listing of services available by location
SLO	IBM Cloud SLO

Related documents

Shared responsibilities for using IBM Cloud products

In IBM Cloud®, as is the case for other cloud service providers, the responsibilities for managing the lifecycle of, operating, and securing products are shared between IBM® and the customer.

The responsibility of completing the following types of tasks on various products can be exclusive to IBM, the customer, or shared between the two. The tasks for each type of product are grouped in the following categories:

- Incident and operations management: Includes tasks such as monitoring, event management, high availability, problem determination, recovery, and full state backup and recovery.
- Change management: Includes tasks such as deployment, configuration, upgrades, patching, configuration changes, and deletion.
- Identity and access management: Includes tasks such as authentication, authorization, access control policies, and approving, granting, and revoking access.
- Security and regulation compliance: Includes tasks such as security controls implementation and compliance certification.
- Disaster recovery: Includes tasks such as providing dependencies on disaster recovery sites, provision disaster recovery environments, data and configuration backup, replicating data and configuration to the disaster recovery environment, and failover on disaster events.

When you're reviewing the following sections, the tables list resources for each category and who manages them. The following list describes what constitutes each type of resource in IBM Cloud.

Data

: Customer-owned content that includes all data that is managed and controlled by the customer. Examples include information that is stored into volumes, files, and databases hosted on IBM Cloud resources and data processed, stored, and logged by the client applications hosted on IBM Cloud. It doesn't include client metadata, the information that is used by IBM to provide services to the customer and support and operate the client account, services, and resources that are always considered to be shared responsibility between client and IBM.

Applications

: Customer-owned software components, such as executables, web applications, middleware, frameworks, libraries, and other software packages that the client developed or acquired by third parties and deployed in IBM Cloud.

Service instance

: An entity that consists of resources that are reserved for a particular service.

Operating systems

: The operating system software and configuration that are deployed in virtual or bare metal servers, such as Linux, Windows, or similar to the ones provided in stock images.

Virtual and bare metal servers

: The virtual or bare metal servers that are ordered and managed through IBM Cloud services.

Virtual storage

: The block, file, or Object Storage buckets ordered and managed through IBM Cloud.

Virtual network

: Network resources such as VLAN, VPC, subnets, or IPs provided by classic infrastructure and VPC services that are ordered and managed through IBM Cloud.

Hypervisor

: The software and configuration that is deployed in physical servers to host and manage the lifecycle of virtual servers.

Physical servers and memory

: The physical compute devices and resources, such as cores, memory, and GPUs used to host the virtual or bare metal servers.

Physical storage

: The physical storage devices and resources, such as disks and storage devices that are used to host the virtual block, file, or Object Storage buckets.

Physical network and devices

: The physical network devices and resources, such as switches, routers, gateways, firewalls, and load balancers that are used to host the virtual network resources.

Facilities and data centers

: The physical data center buildings with power, cooling, and rooms for all the IBM Cloud physical equipment.

Note: IBM Cloud supports the following types of products and the corresponding shared responsibility models. For more information about each specific service, see the documentation for that service.

Infrastructure-as-a-service

Infrastructure-as-a-service (IaaS) products that are managed by IBM are multi-tenant, accessed remotely, hosted on IBM physical infrastructure, created in customer-owned accounts, and have control plane and data plane security that is owned by IBM. Examples of this product type are Virtual Servers and Bare Metal Servers with the related block volumes that are connected to the customer account private subnets. You can find a list of these types of products in the IBM Cloud catalog on the Services tab. Each product is in an infrastructure subcategory within the Compute or VPC infrastructure categories.

Resource	Incident and Operations Management	Change Management	Identity and Access Management	Security and Regulation Compliance	Disaster Recovery
Data	Customer	Customer	Customer	Customer	Customer
Application	Customer	Customer	Customer	Customer	Customer
Operating system	Customer	Customer	Customer	Customer	Customer
Virtual and bare metal servers	Shared	Shared	Shared	Shared	Shared
Virtual storage	Shared	Shared	Shared	Shared	Shared
Virtual network	Shared	Shared	Shared	Shared	Shared
Hypervisor	IBM	IBM	IBM	IBM	IBM
Physical servers and memory	IBM	IBM	Shared	Shared	IBM
Physical storage	IBM	IBM	IBM	IBM	IBM
Physical network and devices	IBM	IBM	IBM	IBM	IBM
Facilities and data centers	IBM	IBM	IBM	IBM	IBM

Shared responsibilities for IaaS products

Note: For areas marked as shared responsibilities, the customer is responsible for all the configurations, and IBM is responsible for all underlying management. For disaster recover, the customer is responsible for creating resources in a secondary region and managing the application and data disaster recovery.

IBM Hybrid Cloud

IBM Hybrid Cloud products are managed by IBM, hosted on IBM-owned physical infrastructure, and located on-premises in customer-owned locations. An example of this product type is Power Virtual Server Private Cloud.

Resource	Incident and Operations Management	Change Management	Identity and Access Management	Security and Regulation Compliance	Disaster Recovery
Data	Customer	Customer	Customer	Customer	Customer

Application	Customer	Customer	Customer	Customer	Customer
Operating system	Customer	Customer	Customer	Customer	Customer
Virtual and bare metal servers	Shared	Shared	Shared	Shared	Shared
Virtual storage	Shared	Shared	Shared	Shared	Shared
Virtual network	Shared	Shared	Shared	Shared	Shared
Hypervisor	IBM	IBM	IBM	IBM	IBM
Physical servers and memory	IBM	IBM	Shared	Shared	IBM
Physical storage	IBM	IBM	IBM	IBM	IBM
Physical network and devices: Top of rack and spine	IBM	IBM	IBM	IBM	IBM
Physical network and devices: Datacenter core	Customer	Customer	Customer	Customer	Customer
Facilities and data centers	Customer	Customer	Customer	Customer	Customer

Shared responsibilities for hybrid cloud IaaS products

Note: For areas marked as shared responsibilities, the customer is responsible for all the configurations, and IBM is responsible for all underlying management. For disaster recover, the customer is responsible for creating resources in a secondary region and managing the application and data disaster recovery.

Managed products

Products that are managed by IBM require customer responsibilities only for the data or applications that customers add to the service. They are multi-tenant, accessed remotely, hosted on IBM virtual resources, created in IBM-owned accounts, and have control plane and data plane security that is owned by IBM. Examples of this product type are IBM Cloud databases or IBM Cloudant database instances. You can find a list of these types of products in the IBM Cloud catalog on the Services tab. However, any products that are listed in an infrastructure subcategory are infrastructure-as-a-service type products.

Resource	Incident and Operations Management	Change Management	Identity and Access Management	Security and Regulation Compliance	Disaster Recovery
Data	Customer	Customer	Customer	Customer	Customer
Application	Customer	Customer	Customer	Customer	Customer
Service instance	IBM	IBM	IBM	IBM	Shared
Virtual and bare metal servers	IBM	IBM	IBM	IBM	IBM
Virtual storage	IBM	IBM	IBM	IBM	IBM
Virtual network	IBM	IBM	IBM	IBM	IBM
Hypervisor	IBM	IBM	IBM	IBM	IBM
Physical servers and memory	IBM	IBM	IBM	IBM	IBM

Physical storage	IBM	IBM	IBM	IBM	IBM
Physical network and devices	IBM	IBM	IBM	IBM	IBM
Facilities and data centers	IBM	IBM	IBM	IBM	IBM

Shared responsibilities for fully-managed products

Note: For disaster recovery, IBM is responsible to ensure that other regions that are not impacted by the disaster are fully operational and will recover the impacted region by the disaster as quickly as possible.

Managed products on customer's resources

Managed products on customer's resources are orchestrated by IBM. They are single-tenant and data plane products. They are accessed locally in customer accounts, data plane hosted on virtual resources in the customer's account, control plane security owned by IBM, and data plane security owned by the customer. IBM Cloud products of this type include IBM Cloud Kubernetes Service on classic infrastructure and Red Hat® OpenShift® on IBM Cloud® on classic infrastructure.

Resource	Incident and Operations Management	Change Management	Identity and Access Management	Security and Regulation Compliance	Disaster Recovery
Data	Customer	Customer	Customer	Customer	Customer
Application	Customer	Customer	Customer	Customer	Customer
Service instance	Shared	Shared	Shared	Shared	Shared
Operating system	Shared	Shared	Shared	Shared	Shared
Virtual and bare metal servers	Shared	Shared	Shared	Shared	Shared
Virtual storage	Shared	Shared	Shared	Shared	Shared
Virtual network	Shared	Shared	Shared	Shared	Shared
Hypervisor	IBM	IBM	IBM	IBM	IBM
Physical servers and memory	IBM	IBM	IBM	IBM	IBM
Physical storage	IBM	IBM	IBM	IBM	IBM
Physical network and devices	IBM	IBM	IBM	IBM	IBM
Facilities and data centers	IBM	IBM	IBM	IBM	IBM

Shared responsibilities for self-managed products

Note: For areas marked as shared responsibilities, the customer is responsible for all the configurations, and IBM is responsible for all underlying management. For disaster recovery, the customer is responsible for creating resources in a secondary region and managing the application and data disaster recovery.

Software packages

Software packages are deployed by IBM as single tenant instances, and they are accessed locally in the customer account. The software instance is hosted on resources in the customer's accounts. The software deployment control plane security is owned by IBM, and the software instance security is owned by the customer.

A generic software deployment control plane manages the lifecycle of deployed software package instances. At a minimum, it manages the deployment, upgrade, and delete actions. As the packages become smarter, the generic control plane might also manage the start, stop, migration, scaling, monitoring, backup, and restore tasks.

You can find a list of software in the IBM Cloud catalog on the Software tab.

Resource	Incident and Operations Management	Change Management	Identity and Access Management	Security and Regulation Compliance	Disaster Recovery
Data	Customer	Customer	Customer	Customer	Customer
Application	Customer	Customer	Customer	Customer	Customer
Software packages	Shared	Shared	Customer	Customer	Shared
Operating system	Shared	Shared	Customer	Customer	Shared
Virtual and bare metal servers	Shared	Shared	Shared	Shared	Shared
Virtual storage	Shared	Shared	Shared	Shared	Shared
Virtual network	Shared	Shared	Shared	Shared	Shared
Hypervisor	IBM	IBM	IBM	IBM	IBM
Physical servers and memory	IBM	IBM	IBM	IBM	IBM
Physical storage	IBM	IBM	IBM	IBM	IBM
Physical network and devices	IBM	IBM	IBM	IBM	IBM
Facilities and data centers	IBM	IBM	IBM	IBM	IBM

Shared responsiblities for software packages

Note: For areas marked as shared responsibilities, the customer is responsible for all the configurations, and IBM is responsible for all underlying management. For disaster recovery, the customer is responsible for creating resources in a secondary region and managing the application and data disaster recovery.

Deployable architectures

Deployable architectures can be comprised of infrastructure-as-a-service, managed products, and software. For more information about the specific responsibilities for you and for IBM when you use a deployable architecture, see Understanding your responsibilities when using deployable architectures.

Public peering

Public peering is conducted across a shared network. Peering requests can be accepted when a mutually agreeable decision to peer based on operational needs exists.

Guidelines and general information

The following are guidelines and some general information for public peering with IBM Cloud®.

- Potential and current peers must have and maintain a valid PeeringDB.com record.
- Use of the Internet Routing Registry (IRR) is required for prefix filtering.
- Peers must supply a contact with 24x7 availability, such as a Network Operations Center.
- Peers must not use any form of gateway of last resort or default route that is directed at IBM Cloud.
- Peers must provide access to a looking glass or traceroute server to facilitate troubleshooting.
- Peer announcements must be consistent across all Interconnect locations.
- IBM Cloud sends and accepts MED attributes by request only.

IBM Cloud reserves the right to suspend peering for an indefinite period if any form of network abuse is detected and verified to be taking place through the peering interconnect. This suspension includes setting a default route to IBM Cloud. IBM Cloud doesn't accept any IPv4 announcements smaller than /24 or IPv6 announcements smaller than /48. An MD5 password is preferred but not required.

For more information, see the PeeringDB record.

Note: If IBM Cloud's operational needs cause a peering arrangement to conflict with the best interest of IBM Cloud, IBM Cloud reserves the right to end the peering agreement with reasonable advance notice.

Public peer exchanges

The following table shows the current public peer exchanges.

Country	City	State or Province	Exchange
Brazil	São Paulo	São Paulo	PTT-SP
Canada	Toronto	Ontario	TorIX
United States	Denver	Colorado	ANY2 Denver
United States	Denver	Colorado	IX Denver
United States	Miami	Florida	FL-IX
United States	Miami	Florida	Equinix Miami
United States	Boston	Massachusetts	MASS-IX
United States	New York	New York	DE-CIX NYC
United States	New York	New York	TIE NYC
United States	New York	New York	NYIIX
United States	Seattle	Washington	SIX
United States	Los Angeles	California	ANY2 Los Angeles
United States	Los Angeles	California	Equinix Los Angeles

United States	Palo Alto	California	Equinix Palo Alto
United States	San Jose	California	Equinix San Jose
United States	Atlanta	Georgia	TIE-ATL
United States	Chicago	Illinois	Equinix Chicago
United States	Dallas	Texas	Equinix Dallas
United States	Dallas	Texas	DE-CIX DAL
United States	Ashburn	Virginia	Equinix Ashburn
United States	Ashburn	Virginia	LINX NoVa

Public peer exchanges for North and South America

Country or region	City	State or Province	Exchange
Australia	Melbourne	Victoria	Equinix Melbourne
Australia	Melbourne	Victoria	IX Australia Victoria
Australia	Melbourne	Victoria	MegaIX Melbourne
Australia	Sydney	New South Wales	Equinix Sydney
Australia	Sydney	New South Wales	IX Australia New South Wales
Australia	Sydney	New South Wales	MegaIX Sydney
Hong Kong S.A.R. of the PRC	New Territories	Hong Kong S.A.R. of the PRC	Equinix Hong Kong
Hong Kong S.A.R. of the PRC	New Territories	Hong Kong S.A.R. of the PRC	HKIX
Hong Kong S.A.R. of the PRC	New Territories	Hong Kong	BBIX Hong Kong
India	Mumbai	Mumbai	AMS-IX Mumbai
India	Mumbai	Mumbai	DE-CIX Mumbai
India	Mumbai	Mumbai	Extreme-IX Mumbai
Japan	Osaka	Japan	BBIX Osaka
Japan	Osaka	Japan	JPIX Osaka
Japan	Osaka	Japan	JPNAP Osaka
Japan	Tokyo	Japan	BBIX Tokyo
Japan	Tokyo	Japan	Equinix Tokyo
Japan	Tokyo	Japan	JPNAP Tokyo
Singapore	Singapore	Singapore	BBIX Singapore

Singapore	Singapore	Singapore	Equinix Singapore
Taiwan	Taipei	Taipei	TPIX-TW

Public peer exchanges for Asia Pacific

Country	City	State or Province	Exchange
Germany	Frankfurt	Hesse	DE-CIX
Germany	Frankfurt	Hesse	NIX.CZ
England	London	England	LINX Juniper LAN
England	London	England	LINX Extreme LAN
France	Paris	Île-de-France	Equinix Paris
Italy	Milan	Italy	MIX-IT
Netherlands	Amsterdam	North Holland	AMS-IX
Netherlands	Amsterdam	North Holland	NL-IX
Spain	Madrid	Madrid	DE-CIX Madrid
Spain	Madrid	Madrid	ESPANIX Madrid Lower LAN
Spain	Madrid	Madrid	ESPANIX Madrid Upper LAN
Spain	Madrid	Madrid	Equinix Madrid

Public peer exchanges for Europe

National language support for IBM Cloud

IBM Cloud® supports nine national languages other than English. However, depending on the release level of the products included in IBM Cloud, the content that is provided might not be translated.

The IBM Cloud console and the Support Center display content according to your browser's language setting. You can switch languages for the console user interface and Support Center by going to your browser's settings page and changing the language.

The IBM Cloud docs are machine translated. You can switch between languages by using the **Language** switcher at the beginning of each page. The language that you select persists between sessions. If it's the first time you're viewing the docs in a particular browser, the docs display the content according to your browser's language setting unless the URL includes a parameter setting the language.

The IBM Cloud AI assistant supports all languages that are supported by IBM Cloud including: Deutsch, English, Español, Français, Italiano, 日本語, 한국어, Português/Brasil, 简体中文, and 繁體中文. To change the language that the AI assistant reponds in, you must ask a question in one of the supported languages first. The browser language setting does not affect generated responses, so simply ask your question in one of the supported languages and get a response back in that same language.

The following table lists the supported national languages and language codes for IBM Cloud.

National Language	Language Code
Brazilian Portuguese	pt_BR
English	en
French	fr
German	de
Japanese	ja
Korean	ko
Italian	it
Spanish	es
Simplified Chinese	zh_CN
Traditional Chinese	zh_TW

Supported national languages and language codes

National language support for support

Support by local technical account managers is available in any of the supported national languages that are listed during local business hours. Our major centers have several bilingual technical support professionals who support local languages as needed.

National language support for sales

IBM has a local presence in most countries, with sales and technical resources that speak the local language. The IBM Cloud website does IP detection to designate the website language, and you can also change your business location in the **Let's talk** button. There, you can find a local phone number and contact information for your country.

The following list displays the website URL and local phone number for a selection of locations that IBM Cloud sales supports across the world.

- United States https://www.ibm.com/cloud / telephone: +18449525683
- Brazil https://www.ibm.com/br-pt/cloud / telephone: +558007071426
- France https://www.ibm.com/fr-fr/cloud / telephone: +33810005810
- Germany https://www.ibm.com/de-de/cloud / telephone: +491805426014
- Japan https://www.ibm.com/jp-ja/cloud / telephone: 81-3-6667-1111

- Korea https://www.ibm.com/kr-ko/cloud / telephone: +82237815156
- Italy https://www.ibm.com/it-it/cloud / telephone: +390270316122
- Spain https://www.ibm.com/es-es/cloud / telephone: +34902022007
- China https://www.ibm.com/cn-zh/cloud / telephone: +864006682350

Glossary terms for IBM Cloud

This glossary provides terms and definitions for IBM Cloud®.

The following cross-references are used in this glossary:

- *See* refers you from a nonpreferred term to the preferred term or from an abbreviation to the spelled-out form.
- *See also* refers you to a related or contrasting term.

A

access control list

A list that statelessly manages inbound and outbound traffic for a subnet through the use of rules. An access control list helps provide security at the subnet level.

access group

A set of users and service IDs organized into a group that is used as the subject of an access policy for assigning all group members the same access.

access policy

A method for granting users, service IDs, and access groups access to account resources. An access policy includes a subject, target, and role.

access token

A value used by the consumer to gain access to the protected resources on behalf of the user, instead of using the user's service provider credentials.

account

A unit within the representation of an organizational structure. It can include resources, such as servers, access controls, and configurations. See also enterprise account.

account group

An organizational unit for accounts within an enterprise. An account group can contain accounts or other account groups. See also enterprise.

accuracy

A measure of the correctness of the decisions and predictions that are made by machine learning models, which are often part of AI systems. See also precision, recall.

accuracy analysis

Analyzing machine learning model scores to determine whether changes are needed to improve accuracy.

action

- A code snippet that can be explicitly invoked, or run in response to an event. See also feed, invoke.
- A task that is performed in the context of a service.

adjudication

An iterative process for resolving annotation conflicts by comparing the annotations added to the same document by different human annotators.

affinity

Two or more container group instances running on the same network node. See also anti-affinity.

agent

A process that performs an action on behalf of a user or other program without user intervention or on a regular schedule, and reports the results back to the user or program.

AI

See artificial intelligence.

AI system

See artificial intelligence system.

allowlist

A list of items, such as usernames, email addresses, or IP addresses, that are granted access to a certain system or function. When an allowlist is used for access control, all entities are denied access, except for those that are included in the allowlist. See also blocklist.

analysis engine

A program that analyzes artifacts, such as documents, and infers information about them, and which implements the UIMA Analysis Engine interface specification. Analysis engines are constructed from building blocks called annotators. An analysis engine can contain a single annotator, which is referred to as a primitive analysis engine, or multiple annotators, which is referred to as an aggregate analysis engine.

annotation

Information about a span of text. For example, an annotation might indicate that a span of text represents a company name.

annotation process manager

A role that is responsible for managing the full annotation lifecycle activities within a workspace. The project manager that is added to a workspace typically performs the activities of an annotation process manager.

annotation set

- In machine-based annotation, a collection of documents that can be used as blind data, training data, or test data.
- In human annotation, a collection of documents that are extracted from the corpus that allow the workload to be shared by multiple human annotators.

anti-affinity

Two or more container group instances that run on different network nodes to ensure higher availability for an app. See also affinity.

API key

A unique code that is used to authenticate and authorize API requests. The code is passed to an API to identify the calling application or user and to track and control how the API is used.

API operation

A unit of a REST API that can be invoked. An API operation comprises an HTTP verb and a URL path that is subordinate to the context root of the API.

API resource

A unit of a REST API that can be invoked. An API resource comprises an HTTP verb and a unique URL path that is subordinate to the context root of the API.

app

A web or mobile device application. See also web application, mobile application.

artifact

An entity that is used or produced by a software or systems development process. Examples of artifacts include designs, requirements, source files, plans, scripts, simulations, models, test plans, and binary executable files. In an HTTP context, artifacts have a URI and are called resources.

artificial intelligence (AI)

The capability to acquire, process, create and apply knowledge in the form of a model to make predictions, recommendations or decisions.

artificial intelligence system (AI system)

A system that can make predictions, recommendations or decisions that influence physical or virtual environments, and whose outputs or behaviors are not necessarily pre-determined by its developer or user. AI systems are typically trained with large quantities of structured or unstructured data, and might be designed to operate with varying levels of autonomy or none, to achieve human-defined objectives.

assembly

An application programming interface that provides rich functionality for interacting with an application. The assembly makes side calls to external services and then transforms and aggregates the response before a response is relayed to the calling application.

asset

Tangible or intangible goods, services, or property represented as an entity that is traded on a blockchain network.

attachment

In Security and Compliance Center, the connection between a profile and scope.

attribute

A characteristic or trait of an entity that describes the entity; for example, the telephone number of an employee is one of the employee attributes.

authentication (AuthN)

The process of validating the identity of a user or server.

AuthN

See authentication.

authorization (AuthZ)

In computer security, the right granted to a user to communicate with or make use of a computer system.

AuthZ

See authorization.

availability zone

A location within a region that IBM Cloud Kubernetes Service runs in.

B

bare metal server

A dedicated, fully-customizable physical server that can be used for virtualization or web hosting.

base image

An image that has no parent image. See also image, parent image.

beta product

A product that IBM makes available solely for evaluation and testing purposes. There are no warranties, SLAs or support provided and beta products are not intended for production use.

bias detection

The process of calculating fairness to metrics to detect when AI models are delivering unfair outcomes based on certain attributes.

billing option

The method by which a client is billed for cloud service usage. Examples include paying in advance, such as with subscriptions, and payment in arrears, as in Pay-As-You-Go accounts.

billing unit

The highest level billing entity within an enterprise. Each account or account group is linked to a billing unit, which manages all associated contracts, invoices, orders, and payments. A billing unit can contain one or more credit pools. See also enterprise, credit pool.

bind

To establish a connection between software components on a network using an agreed-to protocol. In web services, the bind operation occurs when the service requester invokes or initiates an interaction with the service at run time using the binding details in the service description to locate, contact, and invoke the service.

blind data

A set of documents annotated with the ground truth, such as question and answer pairs, semantic annotation, and passage judgment. Blind data is never released or seen by developers and is used to test the system periodically to evaluate performance on unseen data. See also training data, testing data.

block

A unit that contains ordered transactions in a blockchain network. Blocks are immutable and contain the cryptographic hash of the previous block to create a chain of blocks. This iterative process confirms the integrity of the previous block, all the way back to the initial block. See also genesis block.

blocklist

A list of items, such as usernames, email addresses, or IP addresses, that are denied access to a certain system or function. When a blocklist is used for access control, all entities are allowed access, except for those that are included in the blocklist. See also allowlist.

BLU Acceleration

A collection of IBM Db2 technologies designed to work primarily with read-mostly business intelligence query processing. BLU Acceleration consists of four major database design advances: dynamic in-memory columnar processing, actionable compression, parallel vector processing, and data skipping.

blue-green deployment

A deployment technique that enables continuous delivery and minimizes downtime by running two nearly identical production environments called Blue and Green. While one of the environments (for example, Blue) is the live production environment, the other (for example, Green) can be used for final testing and deployment. After the application is deployed in Green, Green becomes the production environment and Blue becomes idle. See also red-black deployment.

boilerplate

A template that includes one application and its associated runtime environment and predefined services for a particular domain.

borderless

Pertaining to an open, non-proprietary development platform that includes public cloud, dedicated cloud, and local cloud deployment models. See also public cloud, local cloud, dedicated cloud.

bucket

A container for storing unstructured data.

buildpack

A collection of scripts that provide framework and runtime support for apps.

business continuity

The capability of a business to withstand outages and to operate mission-critical services normally and without interruption in accordance with predefined service-level agreements.

C

CA

See certificate authority.

certificate authority (CA)

A trusted third-party organization or company that issues the digital certificates. The certificate authority typically verifies the identity of the individuals who are granted the unique certificate. See also Secure Sockets Layer, intermediate certificate, trusted root.

certificate signing request (CSR)

An electronic message that an organization sends to a certificate authority (CA) to obtain a certificate. The request includes a public key and is signed with a private key; the CA returns the certificate after signing with its own private key.

chaincode

Executable code that contains business logic agreed to by a set of organizations on a channel. A chaincode, and the smart contracts it contains, is installed on peers, instantiated on a channel, and governs access to the ledger data that is generated through invocations of the chaincode. See also transaction.

channel

A private subset of a larger blockchain network with specific rules and a separate ledger that only channel members can access.

CLI

See command-line interface.

client

- An entity that acts on behalf of a user by connecting to a peer to communicate with the blockchain.
- A software program or computer that requests services from a server. See also host.

client secret

A piece of information that is used with an application key to verify the identity of an application. An API can be configured to require that client applications supply their application secret with their application key. The application secret functions effectively as a password known only to the application. The application secret is passed by the client using an HTTP query parameter.

cloud computing

A computing platform where users can have access to applications or computing resources, as services, from anywhere through their connected devices. A simplified user interface or application programming interface (API), or both, makes the infrastructure supporting such services transparent to users.

cloud portability

The ability to move applications and services across public or private cloud computing environments, or from different cloud providers.

cloud provider

An organization that provides cloud computing resources.

cloud resource name (CRN)

A globally unique identifier for a specific cloud resource. The value is segmented hierarchically by version, instance, type, location, and scope, separated by colons.

command-line interface (CLI)

A computer interface in which the input and output are text based.

community

A collection of consumer organizations. It is used as a grouping construct when publishing APIs. Communities are used to restrict the visibility and accessibility of APIs.

component

An object that is used to construct a service or resource in a stream or workspace. Examples include a physical server, disk, network switch, power supply, battery back up.

compute

Infrastructure or resources that serve as the basis for building apps in the cloud.

config rule

See configuration rule.

configuration rule (config rule)

A JSON document that defines the configuration of resources and validates the compliance based on security requirements when a resource is created or modified.

confusion matrix

A performance measurement that determines the accuracy between a model's positive and negative predicted outcomes compared to positive and negative actual outcomes.

connection

In data communication, an association established between entities for conveying information.

consensus

The process of participants in a blockchain agreeing to a transaction and validating it through the peer network. Consensus ensures that shared ledgers are exact copies, and lowers the risk of fraudulent transactions since tampering would have to occur across many places at the exact same time.

consumer

A member in a blockchain network that uses the network to invoke transactions against the distributed ledger.

container

A system construct that allows users to simultaneously run separate logical operating system instances. Containers use layers of file systems to minimize image sizes and promote reuse. See also image, registry, layer.

context

Environmental or conditional information that can be used to define when rules that restrict access to resources should be applied. A context can be a set of properties or attributes like client IP addresses, office hours, time of day, or multi-factor authentication.

context-based restriction

An access management method that defines and enforces access for resources based on the network location of the access request.

control

A technical, administrative, or physical safeguard designed to meet a set of defined security and privacy requirements. Controls exist to prevent, detect, or lessen the ability of a threat to exploit a vulnerability.

control assessment

The evaluation of a configuration for compliance with applicable standards.

control library

A collection of similar predefined or custom controls.

control specification

A statement that defines the specific security and privacy requirements that a control must meet.

coreference

A relationship between two words or phrases in which both refer to the same person or thing and one stands as a linguistic antecedent of the other. For example, there is a coreference between the two pronouns in the phrase "She taught herself" but not in the phrase "She taught her". A coreference links two equivalent entities in the same text.

coreference chain

A list of entities that were annotated as coreferences. When a mention is annotated as a coreference, the system creates a coreference chain. The coreference chain provides a way to view all of the mentions in context and verify that all of the occurrences belong together under the same entity type.

corpus

A collection of source documents that are used to train a machine learning model.

credential

Information acquired during authentication that describes a user, group associations, or other security-related identity attributes, and that is used to perform services such as authorization, auditing, or delegation. For example, a user ID and password are credentials that allow access to network and system resources.

credit pool

Within an enterprise billing unit, a consolidation of credit from all sources, including subscriptions and promotions, that is shared among accounts. See also billing unit.

CRN

See cloud resource name.

cryptosystem

A suite of algorithms that are used to implement a specific security service. Examples include RSA and Ed25519.

crypto unit

A single unit that represents a hardware security module and the corresponding software stack that is dedicated to the hardware security module for cryptography.

CSR

See certificate signing request.

curate

To select, collect, preserve, and maintain content relevant to a specific topic. Curation establishes, maintains, and adds value to data; it transforms data into trusted information and knowledge.

custom domain

The customized portion of the URL selected by the user to direct requests to the application. A custom domain makes up part of the route. A custom domain can be a shared domain, a shared subdomain, or a shared domain and host. See also host, subdomain, Uniform Resource Locator, domain, route.

D

daemon

A program that runs unattended to perform continuous or periodic functions, such as network control.

dashboard

A user interface component that provides a comprehensive summary of pertinent information from various sources to the user.

data center (DC)

The physical location of the servers that provide cloud services.

data encryption key

A cryptographic key used to encrypt data that is stored in an application.

data store

A place, such as a database system, file, or directory, where data is stored.

DC

See data center.

dedicated cloud

A private cloud computing environment that provides infrastructure with single-tenant hardware. See also borderless.

deep learning

A computational model that uses multiple layers of interconnected nodes, which are organized into hierarchical layers, to transform input data (first layer) through a series of computations to produce an output (final layer). Deep learning is inspired by the structure and function of the human brain.

deployable architecture

Cloud automation for deploying a common architectural pattern that combines one or more cloud resources that is designed for easy deployment, scalability, and modularity. See also module.

deployment

A process that retrieves the output of a build, packages the output with configuration properties, and installs the package in a pre-defined location so that it can be tested or run. See also stage.

DevOps

A software methodology that integrates application development and IT operations so that teams can deliver code faster to production and iterate continuously based on market feedback.

DevSecOps

A methodology that integrates security practices with the software development and operations lifecycle. The goal of the merge is to prioritize the balance of development speed and security.

dictionary

A collection of words that can be used to pre-annotate documents. A new annotation is created for each word in the document text that matches a term in the dictionary. A machine learning model can be configured with one or more independent dictionaries, which are typically domain-specific, such a dictionary for pharmaceuticals and a dictionary for wealth management. See also lemma, surface form.

dictionary pre-annotator

A component that identifies mentions in text that match a specific set of words. By using domain-specific terminology to pre-annotate text, dictionary pre-annotators can accelerate a human annotator's ability to prepare a set of ground truth documents.

disaster recovery (DR)

The ability of a service or workload to recover from rare, major incidents and wide-scale failures, such as service disruption. This includes a physical

disaster that affects an entire region, corruption of a database, or the loss of a service contributing to a workload. The impact exceeds the ability of the high availability design to handle it. See also high availability, recovery time objective.

Dockerfile

A text file that contains instructions to build a Docker image.

document set

A collection of documents. Documents that are imported together become a document set. Annotated documents that are grouped together for training purposes are generated as document sets.

domain

Part of a naming hierarchy that specifies the route. For example, example.com. In IBM Cloud, domains are associated with orgs. Domain objects are not directly bound to apps. See also host, subdomain, Uniform Resource Locator, custom domain, organization, route.

DR

See disaster recovery.

dynamic secret

A unique value, such as a password or an API key, that is created dynamically and leased to an application that requires access to a protected resource. After a dynamic secret reaches the end of its lease, access to the protected resource is revoked and the secret is deleted automatically. See also secret, secrets engine.

E

endorse

To validate a chaincode transaction that was made by another member of a blockchain network.

endorsement

A collection of digital signatures from endorsing peers that establish that a transaction satisfies an endorsement policy.

endorsement policy

A policy that defines the peer nodes on a channel that must execute transactions that are attached to a specific chaincode application, and the required combination of endorsements. For example, a policy could require that a transaction be endorsed by a minimum number of endorsing peers, a minimum percentage of endorsing peers, or by all endorsing peers that are assigned to a specific chaincode application.

endpoint

The address of an API or service in an environment. An API exposes an endpoint and at the same time invokes the endpoints of other services. See also route.

enterprise

A hierarchical structure of accounts with centralized account and billing management in a cloud environment. See also account group, billing unit.

enterprise account

A unique account within an enterprise that manages users and access for account organization and enterprise billing. See also account.

entitlement

In software licensing, the maximum allowed allocation of capacity as determined by a license agreement.

entity

- A set of details that are held about a real-world object such as a person, location, or bank account. An entity is a kind of item.

- A person, object, or concept about which information is stored.
- A mention that is annotated by an entity type.

entity type

The type of entity that a mention represents without consideration for context. For example, the mention IBM might be annotated by the entity type ORGANIZATION. In an entity-relationship model, an entity type is the thing that is being modeled or the thing that a mention refers to, such as the name of a person or place. Different entity types have different sets of attributes such as "surname" or "home town", and are connected through relationships like "lives in". An entity type exists independently and can be uniquely identified.

envelope encryption

The process of encrypting data with a data encryption key and then encrypting the key with a root key that can be fully managed.

evidence

The collected raw data to support an assessment for auditing purposes.

experimental product

A product that IBM makes available solely for evaluation and testing purposes, and might be unstable or not compatible with previous versions. An experimental product can be discontinued with short notice. There are no warranties, SLAs or support provided, and experimental products are not intended for production use.

explainability

The ability of human users to trace, audit, and understand predictions that are made in applications that use AI systems.

F

F1 score

A measure of a test's accuracy that considers both precision and recall to compute the score. The F1 score can be interpreted as a weighted average of the precision and recall values. An F1 score reaches its best value at 1 and worst value at 0.

false negative

An answer or annotation that is correct, but was predicted to be incorrect.

false positive

An answer or annotation that is incorrect, but was predicted to be correct.

feature

A data member or attribute of a type.

feature code

A code that is applied to free accounts to unlock extra product resources and capabilities.

Federal Risk and Authorization Management Program (FedRAMP)

A United States government program that provides a standardized, risk-based approach for the adoption and use of cloud services by the US federal government. FedRAMP empowers agencies to use modern cloud technologies with an emphasis on security and protection of federal information, and helps accelerate the adoption of secure cloud solutions.

federate

To merge two or more entities. For example, a company's registered domain could be federated with an IBMid.

FedRAMP

See Federal Risk and Authorization Management Program.

feed

A piece of code that configures an external event source to fire trigger events. See also action.

file share

In the IBM Cloud environment, a persistent storage system where users store and share files. In IBM Cloud Kubernetes Service, users can mount Docker volumes on file shares.

fine tuning

The process of adapting a pre-trained model to perform a specific task by conducting additional training. Fine tuning may involve (1) updating the model's existing parameters, known as full fine tuning, or (2) updating a subset of the model's existing parameters or adding new parameters to the model and training them while freezing the model's existing parameters, known as parameter-efficient fine tuning.

fire

To activate a trigger.

Fleiss Kappa score

A measure of how consistently the same annotation was applied by multiple human annotators across overlapping documents. The Fleiss Kappa score reaches its best value at 1 and worst value at 0.

floating IP address

A public, routable IP address that makes use of 1-to-1 network address translation (NAT) so that a server can communicate with the public internet and private subnet within a cloud environment. Floating IP addresses are associated to an instance, for example, a virtual server instance, a load balancer, or a VPN gateway, by means of a virtual network interface card (vNIC).

foundation model

An AI model that can be adapted to a wide range of downstream tasks. Foundation models are typically large-scale generative models that are trained on unlabeled data using self-supervision. As large scale models, foundation models can include billions of parameters.

framework

An architecture for an application that provides a standard structure for an application, and general, extensible functionality. A framework enables and simplifies consistent implementation of complex technologies for application development.

functional identifier

An ID created through a federated identity provider that represents a program, application, or service that is assigned the minimum level of access required to complete the function for which it is created.

G

GA

See general availability.

GB-hour

The cumulative amount of memory (in gigabytes) that is running for all application instances for a particular buildpack per hour.

gen AI

See generative AI.

general availability (GA)

Date when a product is widely available for sale and delivery to customers or channels, usually across multiple geographies.

generative AI (gen AI)

A class of AI algorithms that can produce various types of content including text, source code, imagery, audio, and synthetic data.

genesis block

The configuration block that initializes a blockchain network or channel, and also serves as the first block on a chain. See also block.

globally unique identifier (GUID)

An algorithmically determined number that uniquely identifies an entity within a system.

gossip

A method of sharing network information among peers in which each peer forwards messages to a random selection of the current peers in the network.

Gossip Data Dissemination Protocol

A protocol for secure, reliable, and scalable communication of information in an network by passing messages among peers.

ground truth

The set of vetted data, consisting of annotations added by human annotators, that is used to adapt a machine learning model to a particular domain. Ground truth is used to train machine learning models, measure model performance (precision and recall), and calculate headroom to decide where to focus development efforts for improving performance. Accuracy of ground truth is essential since inaccuracies in the ground truth will correlate to inaccuracies in the components that use it.

GUID

See globally unique identifier.

H

HA

See high availability.

hardware security module (HSM)

A physical appliance that provides on-demand encryption, key management, and key storage as a managed service.

hardware virtual machine mode (HVM)

Hardware-assisted full virtualization. A virtual machine uses resources from the host computer to operate as a complete hardware environment. The host operating system is unaware of the virtual client.

Hash-Based Message Authentication Code (HMAC)

A cryptographic code that uses a cryptic hash function and a secret key.

headroom analysis

The process of determining how much improvement in accuracy, precision, or recall can be expected by addressing some class of problems that are identified while performing accuracy analysis.

health check

A process that monitors system resources and conditions to determine whether the system is running efficiently. The health check can be configured to report potential problems and to display warnings and fail levels before the integrity of the system is compromised.

heavy API call

A client operation that writes, deletes, or inserts data. Heavy API calls consume more resources than light API calls because they are affecting the data.

See also light API call.

high availability (HA)

The ability of a service or workload to withstand failures and continue providing processing capability according to some predefined service level. For services, availability is defined in the Service Level Agreement. Availability includes both planned and unplanned events, such as maintenance, failures, and disasters. See also disaster recovery.

HMAC

See Hash-Based Message Authentication Code.

host

A virtual or physical device, not including a router, that exists on a network. See also subdomain, client, Uniform Resource Locator, custom domain, domain, route.

HSM

See hardware security module.

HTTP method

An action that is used by the Hypertext Transfer Protocol. HTTP methods include GET, POST, and PUT.

HTTPS

See Hypertext Transfer Protocol Secure.

human annotator

A subject matter expert who reviews, modifies, and augments the results of pre-annotation by identifying mentions, entity type relationships, and mention coreferences. By examining text in context, a human annotator helps determine ground truth and improve the accuracy of the machine learning model.

HVM

See hardware virtual machine mode.

hybrid cloud

A cloud computing environment that consists of multiple public and private resources.

Hyperledger fabric

The implementation of the Linux Hyperledger project. See also Linux Hyperledger project.

Hypertext Transfer Protocol Secure (HTTPS)

An Internet protocol that is used by web servers and web browsers to transfer and display hypermedia documents securely across the Internet.

I

IaaS

See infrastructure as a service.

IAM

See identity and access management.

IBM Cloud

An open-standards, cloud-based platform for building, managing, and running apps of all types, such as web, mobile, big data, and smart devices. Capabilities include Java, mobile back-end development, and application monitoring, as well as features from ecosystem partners and open source—all

provided as-a-service in the cloud.

identity and access management (IAM)

The process of controlling access of authorized users to data and applications, while helping companies comply with various regulatory requirements.

identity provider (IdP)

A service that creates, maintains, and manages identity information, account authentication and credential management services for one or more systems. For example, a user registry for one or more hosts; or a built-in user registry for an application.

IdP

See identity provider.

image

A file system and its execution parameters that are used within a container runtime to create a container. The file system consists of a series of layers, combined at runtime, that are created as the image is built by successive updates. The image does not retain state as the container executes. See also container, registry, namespace, layer, base image, parent image, private image repository.

imprint mode

An operational mode in which crypto units are assigned to a user.

infrastructure as a service (IaaS)

The delivery of a computer infrastructure, including server functionality, networking functionality, data center functionality, and storage functionality as an outsourced service.

input/output operations per second (IOPS)

A standard computing benchmark used to determine the best configuration settings for servers.

instance

An entity that consists of resources that are reserved for a particular application or a service.

instantiate

The process of starting and initializing a smart contract on a blockchain channel. After instantiation, peers on the channel that have the smart contract installed can accept smart contract invocations.

inter-annotator agreement

A measure of how similarly a document in two or more document sets is annotated.

intermediate certificate

A subordinate certificate that is issued by the trusted root certificate authority (CA) specifically to issue end-entity server certificates. The result is a certificate chain that begins at the trusted root CA, passes through the intermediate certificate, and ends with the SSL certificate issued to the organization. See also trusted root, certificate authority.

Internet of Things (IoT)

The global network of endpoints that can capture or generate data. For example, a smartphone, smart watch and back-end server might all communicate with each other, sending data back and forth, or even to additional devices within the network.

Internet Small Computer System Interface (iSCSI)

An IP-based standard for linking data storage devices over a network and transferring data by carrying SCSI commands over IP networks.

inventory

A dynamic set of application and environment data from ingestion jobs or uploaded software bill of materials (SBOM) data.

invoke

To activate an action. See also action.

IOPS

See input/output operations per second.

IoT

See Internet of Things.

iSCSI

See Internet Small Computer System Interface.

J

JAR file

A Java archive file.

JavaScript Object Notation (JSON)

A lightweight data-interchange format that is based on the object-literal notation of JavaScript. JSON is programming-language neutral but uses conventions from various languages.

job

A group of executable tasks and operations that can be run as a batch. As a result, a historical record of the actions performed is created.

JSON

See JavaScript Object Notation.

K

knowledge graph

A model that consolidates typed entities, their relationships, their properties, and hierarchical taxonomies to represent an organization of concepts for a given domain. After the knowledge graph store is loaded with inputs from structured and unstructured data sources, users and applications can access the knowledge graph to explore key elements of knowledge for a specific domain, explore interactions, and discover additional relationships.

L

label

A user-defined identifier attached to a grouping of resources that are contained in an instance. Labels are visible only at an instance.

large language model (LLM)

A language model with a large number of parameters, trained on a large quantity of text.

layer

A changed version of a parent image. Images consist of layers, where the changed version is layered on top of the parent image to create the new image. See also image, container.

LBaaS

See load balancer as a service.

LDAP

See Lightweight Directory Access Protocol.

lemma

The normalized or canonical form of a word. Typically, the lemma is the underived and uninflected form of a noun or a verb. For example, the lemma of the terms 'organizing' and 'organized' is 'organize' See also surface form, dictionary.

light API call

A client operation that only reads data. Light API calls use fewer resources than heavy API calls because they are performing a single function. See also heavy API call.

Lightweight Directory Access Protocol (LDAP)

An open protocol that uses TCP/IP to provide access to directories that support an X.500 model and that does not incur the resource requirements of the more complex X.500 Directory Access Protocol (DAP). For example, LDAP can be used to locate people, organizations, and other resources in an Internet or intranet directory.

Linux Hyperledger project

An open source, collaborative effort to advance blockchain technology by identifying and addressing important features for a cross-industry open standard for distributed ledgers that can transform the way business transactions are conducted globally. Hyperledger serves as the foundation code for the IBM Blockchain products, services, and solutions. See also Hyperledger fabric.

LLM

See large language model.

load balancer

Software or hardware that distributes workload across a set of servers to ensure that servers are not overloaded. The load balancer also directs users to another server if the initial server fails.

load balancer as a service (LBaaS)

A service that provides the ability to distribute traffic among instances in a virtual private cloud.

local cloud

A cloud computing environment within the client's data center. The local cloud is on-premises, providing improved latency and security. See also borderless.

logical unit number (LUN)

In the Small Computer System Interface (SCSI) standard, a unique identifier used to differentiate devices, each of which is a logical unit (LU).

LoopBack data source

A JavaScript object that represents a back-end service such as a database, REST API (to be consumed), or SOAP web service. Data sources are backed by connectors that then communicate directly with the database or other back-end services.

LoopBack model

A model that provides a remote (REST) API that clients use to perform operations and interact with backend systems. The model consists of application data, validation rules, data access capabilities, and business logic. Every LoopBack application by default has a set of built-in models: user, application, email, and several models for access control.

LUN

See logical unit number.

M

machine learning (ML)

A branch of artificial intelligence (AI) and computer science that focuses on the use of data and algorithms to imitate the way that humans learn, gradually improving the accuracy of AI models.

machine learning annotator

See machine learning model.

machine learning model

An AI model that is trained on a a set of data to develop algorithms that it can use to analyze and learn from new data.

Managed Service Provider (MSP)

An IBM Business Partner that provides IT services on a contractual basis to maintain clients' computers, networks or software. They manage services on-site at the clients' data center, remotely in the clients' data center, or in a third-party data center.

master key

An encryption key that is used to protect a crypto unit. The master key provides full control of the hardware security module and ownership of the root of trust that encrypts the chain keys, including the root key and standard key.

MBaaS

See mobile backend as a service.

member

A participant that is enrolled in a blockchain network. A member can be as large as a multi-national corporation or as small as an individual.

membership service provider

A component that defines the organizations that can transact on a blockchain network. The membership service provider turns identities into permissions by defining them as administrators of an organization. It also contains the public certificate for the root of trust of the organization. context: Typically, a single membership service provider is used to represent a single blockchain organization.

mention

A span of text that is considered relevant ina domain data. For example, in a type system about automotive vehicles, occurrences of terms such as "airbag", "Ford Explorer", and "child restraint system" might be relevant mentions.

microservice

A set of small, independent architectural components, each with a single purpose, that communicate over a common lightweight API.

ML

See machine learning.

MLOps

A methodology that takes a machine learning model from development to production.

mobile app

See mobile application.

mobile application (mobile app)

An application that has been designed for a mobile platform. Similar to web applications, mobile apps provide some function beyond static display of information, for example, allowing the user to filter news in near real-time. See also app.

mobile backend as a service (MBaaS)

A computing model that connects mobile applications to cloud computing services and provides features such as user management, push notifications, and integration with social networks through a unified API and SDK.

mobile cloud

An infrastructure in which the storage and processing of data for applications is offloaded from a mobile device into the cloud. With mobile cloud computing, applications are not limited to a specific carrier, but are accessed through the Web.

ModelOps

A methodology for managing the full lifecycle of an AI model, including training, deployment, scoring, evaluation, retraining, and updating.

module

A stand-alone unit of automation code that can be reused and shared in a larger system. See also deployable architecture.

MSP

See Managed Service Provider.

multizone region (MZR)

A region that is spread across physical locations in multiple zones to increase fault tolerance. See also zone.

MZR

See multizone region.

N

named entity

A concept in a domain that falls in to a well defined category, such as names of organizations, locations, authors, or diseases.

namespace

A collection of repositories that store images in a registry. A namespace is associated with an IBM Cloud account, which can include multiple namespaces. See also image, private image repository.

NAT

See network address translation.

natural language processing (NLP)

A field of artificial intelligence and linguistics that studies the problems inherent in the processing and manipulation of natural language, with an aim to increase the ability of computers to understand human languages.

network address translation (NAT)

An addressing method that is used to enable one IP address to communicate with several other IP addresses, such as those on a private subnet, by means of a lookup table. Network address translation has two main types: 1-to-1 and many-to-1.

Network File System (NFS)

A protocol that allows a computer to access files over a network as if they were on its local disks.

NFS

See Network File System.

NLP

See natural language processing.

node

- A member of a cluster that runs as a container or virtual machine on shared hardware.
- A physical component of the blockchain network infrastructure, namely a peer, Certificate Authority (CA), or ordering node.

O

OAuth

An HTTP-based authorization protocol that gives applications scoped access to a protected resource on behalf of the resource owner, by creating an approval interaction between the resource owner, client, and resource server.

OCI container image

A container image that is compliant with the OCI Image Format Specification

on-prem

See on-premises.

on-premises (on-prem)

Pertaining to software that is installed and run on the local computers of a user or organization.

ontology

An explicit formal specification of the representation of the objects, concepts, and other entities that can exist in some area of interest and the relationships among them.

orderer node

See ordering node.

ordering node

A node that collects transactions from network members, orders the transactions and bundles them into blocks.

ordering service

A service that provides a shared communication channel to clients and peers for the broadcast of messages that contain transactions.

org

See organization.

organization (org)

- The grouping methodology for users in IBM Cloud. Orgs are used to manage quotas. Users in an org share memory and service instance quotas. See also domain, space.
- The entity that owns APIs or applications that use APIs. A provider organization owns APIs and associated plans, and can additionally own applications. A consumer organization owns only applications. An organization has at least one owner. An organization can be a project team, department, or division.

origin server

A server that processes and responds to incoming requests from clients, and is typically used with a caching server.

out-of-band

Pertaining to data transferred through a stream that is independent from the main data stream.

P

PaaS

See platform as a service.

paravirtualized mode

A lightweight virtualization technique. While in paravirtualized mode, a virtual machine does not require virtualization extensions from the host computer, thus allowing virtualization on hardware systems that do not support hardware-assisted virtualization.

parent image

An image that provides a base for another image. For example, Ubuntu Linux is the parent image of the IBM Liberty image. See also image, base image.

part of speech (POS)

A grammatical category, such as noun or verb, based on the syntactic function of a lexical item.

path

The route through which users access REST APIs. A path consists of one or more HTTP operations such as GET or POST.

payload logging

The capture of payload data and deployment output to monitor ongoing health of AI in business applications.

payment method

The method by which a client pays an invoice, such as credit card, check, or wire transfer.

PEAR

See processing engine archive.

peer

A blockchain network entity that maintains a ledger with an installed chaincode for performing read and write operations on the ledger. A peer is owned and maintained by an organization.

performance

The measurement of a Watson system in terms of accuracy, precision, and recall, for example, when answering questions, discovering relationships, or annotating text.

plan

The packaging construct by which APIs are made available to consumers. A plan makes available a collection of resources or operations from one or more APIs, and is published to communities of application developers.

platform as a service (PaaS)

The delivery of a computing platform, including applications, optimized middleware, development tools, and Java and Web 2.0 runtime environments, in a cloud-based environment.

pod

A group of containers that are running on a Kubernetes cluster. A pod is a runnable unit of work, which can be a either a stand-alone application or a microservice.

point of presence (PoP)

A physical location that stores servers and routers in a network cloud.

policy

A piece of configuration that controls some aspect of processing in the gateway during the handling of an API invocation. Policies are the building blocks of assembly flows and provide the means to configure capability, such as security, logging, caching, routing of requests to target services, and transformation of data from one format to another. Policies can be configured in the context of an API or in the context of a plan.

PoP

See point of presence.

POS

See part of speech.

pre-annotation

The process of annotating a set of documents prior to human annotation. Documents can be pre-annotated by using a rule-based model, a machine-learning model, IBM Watson Natural Language Understanding, or a dictionary. Pre-annotation can help human annotators more quickly prepare a set of ground truth documents.

precision

A measurement that specifies the proportion of results that are relevant. Precision, which is a positive predictive value, is determined by the number of correct positive results divided by the number of all positive results. Accuracy is best measured by using both precision and recall. See also recall, accuracy.

private cloud

A cloud computing environment in which access is limited to members of an enterprise and partner networks. See also public cloud.

private image repository

The combination of an organization's IBM Cloud registry and its namespace. The private image repository is used when referencing an image in a command. See also image, namespace.

private key

An algorithmic pattern used to encrypt messages that only the corresponding public key can decrypt. The private key is also used to decrypt messages that were encrypted by the corresponding public key. The private key is kept on the user system and is protected by a password.

private resource

An entry that is visible only to account owners and their included accounts. When resources are created, they are private by default. See also public resource.

private service

A service that is visible only to members of a specified IBM Cloud organization.

processing engine archive (PEAR)

A .pear archive file that includes an Unstructured Information Management Architecture (UIMA) analysis engine and all of the resources that are required to use it for custom analysis.

profile

- A specification of a resource's capacities and capabilities. Different profiles are optimized for different workloads and use cases. A resource's pricing model might depend on its profile.
- In Security and Compliance Center, a grouping of similar controls that can be used to evaluate resource configuration for compliance.A profile can be customized to fit specific use cases by editing the default parameters. See also rule, scope.

project

A collection of artifacts that define and manage resources and Infrastructure as Code deployments.

promo code

A code used to apply limited time credits to Pay-As-You-Go and Subscription accounts.

prompt

Data, such as text or an image, that prepares, instructs, or conditions a foundation model's output.

prompt engineering

The process of designing natural language prompts for a language model to perform a specific task.

prompting

The process of providing input to a foundation model to induce it to produce output.

prompt tuning

An efficient, low-cost way of adapting a pre-trained model to new tasks without retraining the model or updating its weights. Prompt tuning involves learning a small number of new parameters that are appended to a model’s prompt, while freezing the model’s existing parameters.

proxy

An application programming interface that forwards requests to a user-defined backend resource and relays responses back to the calling application.

public cloud

A cloud computing environment in which access to standardized resources, such as infrastructure, multi-tenant hardware, and services, is available to subscribers on a pay-per-use basis. See also private cloud, borderless.

public gateway

The connection of a subnet, with all virtual server instances attached, to the internet. A public gateway uses a many-to-1 network address translation (NAT), which means that thousands of virtual server instances with private addresses can use one public IP address to talk to the public internet.

public resource

An entry that is visible to everyone in the IBM Cloud catalog. Public resources can be built by any provider (IBM or third party providers). See also private resource.

publish

The process of moving an application or product from staging so that the plans and APIs included within it are available for application developers to access and use.

push

To send information from a server to a client. When a server pushes content, it is the server that initiates the transaction, not a request from the client.

push notification

An alert indicating a change or update on a mobile app icon.

Q

quota

The number of resources that can be consumed at an account or service instance level.

R

Raft

A crash fault tolerant ordering service implementation based on the etcd library of the Raft protocol. Raft follows a leader-and-follower model, where a leader node is elected per channel, and its decisions are replicated by the followers.

RAG

See retrieval augmented generation.

read-mostly

Pertaining to data that changes dynamically.

recall

A measurement that specifies the percentage of relevant results returned, out of all available relevant results. Recall, which is a measure of sensitivity, is determined by the number of correct positive results divided by the number of positive results that should have been returned. Accuracy is best measured by using both precision and recall. See also precision, accuracy.

recovery point objective (RPO)

In disaster recovery planning, the time at which data is restored measured in time (seconds, minutes, hours) starting at the recovered instance and ending at the point of disaster.

recovery time objective (RTO)

In disaster recovery planning, the duration of time for a business process to be restored after a disaster. See also disaster recovery.

red-black deployment

A deployment technique that drives continuous delivery by enabling synchronized test, development, and deployment. Initially, development is done on an inactive environment (black) while the active environment continues to take traffic (red). Once deployment starts, both environments go live (red-red) until routing is disabled on the formerly active, previous version environment, then subsequently removed (black) while the new environment serves as the only active environment. See also blue-green deployment.

reference architecture

An opinionated pattern of technologies that work together that includes a summary, an architecture diagram, and a list of modules.

region

An independent geographic territory that consists of one or more zones.

registry

A storage and distribution service that contains public or private images that are used to create containers. See also image, container.

relation

Typically a verb that reflects how entities are related to one another. For example, "lives in" is a relation between a person and a town. A relation links two different entities in the same sentence.

relation type

A binary, unidirectional relationship between two entities. For example, Mary employedBy IBM is a valid relationship; IBM employedBy Mary is not.

Representational State Transfer (REST)

A software architectural style that guides the design and development of the architecture for the web. REST defines a set of constraints for the architecture of Internet-scale distributed hypermedia systems, such as the web.

resource

A physical or logical instance that can be provisioned or reserved. Examples of resources can include storage, processors, memory, databases, clusters, and VMs.

resource group

The environment, and constraints, in which contained resource instances adhere to. A user can be associated with a resource group to enable collaboration.

REST

See Representational State Transfer.

retrieval augmented generation (RAG)

A technique in which a large language model is augmented with knowledge from external sources to generate text. In the retrieval step, relevant documents from an external source are identified from the user's query. In the generation step, portions of those documents are included in the LLM prompt to generate a response grounded in the retrieved documents.

role

- An attribute that provides a context-sensitive meaning of a mention. For example, in the phrase "I went to IBM today", IBM is the mention, Organization is the entity type, and Facility is the role of the entity type.
- A set of permissions or access rights.

root key

A symmetric wrapping key that is used for encrypting and decrypting other keys that are stored in a data service.

route

The URL that is used to direct requests to an application. A route is made up of an optional host (or subdomain) and a domain that are specified when an application is pushed. For example, in the route myapp.example.com, myapp is the host and example.com is the domain. A route can be associated with one or more applications. Unless a custom domain is specified, IBM Cloud uses a default shared domain in the route to an application. See also host, subdomain, Uniform Resource Locator, custom domain, endpoint, domain.

RPO

See recovery point objective.

RTO

See recovery time objective.

rule

- The set of allowed contexts that are associated with an IBM Cloud resource. The defined contexts dictate a user's or service's access to that resource.
- A set of conditional statements that are used to determine whether a build can be promoted. See also profile, scope.

rule set

A set of rules that define patterns for annotating text. If a pattern applies, then the actions of the rule are performed on the matched annotations. A rule typically specifies the condition that must match, an optional quantifier, a list of additional constraints that the matched text must fulfill, and the actions to be taken when a match occurs, such as creating a new annotation or modifying an existing annotation.

runtime

The set of resources used to run the application.

S

SaaS

See software as a service.

sandbox catalog

A catalog in which approvals for publishing and lifecycle actions are bypassed so that it can be used for testing APIs under development.

scale

To increase platform (or system) capacity by adding more application or service instances

SCM

See source control management.

scope

A grouping of resources that can be validated or evaluated for security and compliance. See also rule, profile.

secret

Sensitive information, such as a password or an API key, that is used by an application to access a protected resource. See also dynamic secret.

secret group

The environment and constraints that contained secrets in an instance must adhere to. A user can be associated with a secret group to enable access and collaboration.

secrets engine

A component that serves as a back end for a specific type of secret, such as a password or an API key, within a secrets management service. Depending on its type, a secrets engine can store data, generate secrets on demand, and more. See also dynamic secret.

Secure Shell (SSH)

A network protocol for secure data exchange between two networked devices. The client can use public-key and private-key authentication, or password authentication, to access the remote server.

Secure Sockets Layer (SSL)

A security protocol that provides communication privacy. With SSL, client/server applications can communicate in a way that is designed to prevent eavesdropping, tampering, and message forgery. See also certificate authority.

security definition

A specification of the settings for a particular aspect of API security; for example, the user registry that authenticates access to the API.

security group

A resource that provides rules to filter IP traffic to resources in a virtual private cloud. Rules are stateful, such that packets in response to allowed packets are automatically permitted.

select availability

A production-ready product that is available for sale and accessible to select customers.

service

An IBM Cloud offering that supports the creation, deletion, and management of resources and provides a wide range of services like compute, storage, analytics, and AI to handle business workloads.

service binding

An association between an application and a service instance that contains the application credentials that are used to communicate with the service

instance.

service endpoint

The physical address of a service that implements one or more interfaces.

service ID

An identity that authenticates a service or an application to a cloud environment and other services. A service ID can be assigned access policies and used to enable an application that is deployed to a cloud environment access to cloud services.

session

The period of time after an app is started on a mobile device and the quality assurance product is notified to begin collecting app behavior, issues, and problems.

signature key

An encryption key that is used by the crypto unit administrator to sign commands that are issued to the crypto unit.

signCert

A certificate that any entity, such as an organization or admin, attaches to their proposals or proposal responses. These signCerts are unique to an entity and are checked by the ordering service to make sure they match the signCert on file for that entity.

single-campus multizone region (single-campus MZR)

A region that consists of multiple zones that are located within a single building or campus. Dependencies such as power, cooling, networking, and physical security might be shared but are designed to provide a high degree of fault independence. See also zone.

single-campus MZR

See single-campus multizone region.

single-page application (SPA)

An application that works inside a browser and does not require page reloading during use.

single sign-on (SSO)

An authentication process in which a user can access more than one system or application by entering a single user ID and password.

smart contract

A set of business terms that are embedded into a blockchain and executed with transactions. A smart contract can also include a digital representation of a set of business rules and defines conditions under which transfers occur. A smart contract is implemented using chaincode.

SoE

See system of engagement.

software as a service (SaaS)

A model of software deployment whereby software including business processes, enterprise applications, and collaboration tools, are provided as a service to customers through the cloud.

SOLO

A consensus plugin implementation for Hyperledger Fabric that results in a single ordering service node in the blockchain network.

SOR

See system of record.

source control management (SCM)

An aspect of software configuration management that involves managing changes to collections of files.

SPA

See single-page application.

space

A sub-group within an IBM Cloud org. Users who are members of an org are given access to one or more of its spaces, with permissions associated with a particular role (such as developer, manager, or auditor). Any member of the space can view apps, but only members with the developer role can create apps and add service instances to the space. Apps and service instances are associated with spaces. See also organization.

SSH

See Secure Shell.

SSL

See Secure Sockets Layer.

SSO

See single sign-on.

stage

To deploy an application, service, or instance to a pre-defined location for running or testing before deployment to a production environment. See also deployment.

stanza

A section of a software package that defines either a specific action to be performed on that the software package or a set of conditions under which actions are to be performed on the software package. The complete software package is a stanza that contains a hierarchy of many different stanzas.

state database

A database that contains the current values for all keys on a blockchain ledger for a channel.

subdomain

A domain that makes up a part of a larger domain. See also host, Uniform Resource Locator, custom domain, domain, route.

subject

The user, service ID, or access group that is granted access by an access policy.

subnet

See subnetwork.

subnetwork (subnet)

A network that is divided into smaller independent subgroups, which still are interconnected.

subscription code

A code used to apply credits to an account after a subscription is purchased.

subtype

A type that extends or implements another type; the supertype.

surface form

The form of a word or multiword unit as it is found in the corpus. For example, some surface forms of the lemma 'organize' are the terms 'organizing' and 'organized'. See also lemma, dictionary.

system of engagement (SoE)

An information technology (IT) system that incorporates technologies that encourage user interaction through email, collaboration systems, and networking. A system of engagement often uses cloud technologies to extend the usefulness of systems of record. See also system of record.

system of record (SOR)

An information storage system (such as a database or application) that stores business records and automates standard processes. See also system of engagement.

T

tag

A user-defined identifier attached to a grouping of resources that are contained in an account. Tags are visible account-wide.

target

The resource or set of resources to provide a subject access to in an access policy. The set of resources is defined by one or more attributes. For example, a target could be all resources in a resource group, all resources of a certain resource type, or the resource with a certain resource ID.

template

A predefined structure for an artifact.

testing data

A set of annotated documents that can be used to evaluate system metrics after ingestion and training. See also training data, blind data.

third-party

Pertaining to a product or service that is provided by a company other than IBM.

tile

A visual representation of a running application that provides status on a dashboard.

train

To set up a Watson instance with components that enable the system to function in a particular domain (for example: corpus content, training data that generates machine learning models, programmatic algorithms, annotators, or other ground truth components) and then making improvements and updates to these components based on accuracy analysis.

training data

A collection of data that is used to train machine learning models. See also testing data, blind data.

transaction

The mechanism that participants on the blockchain network use to interact with assets. A transaction either creates new chaincode or invokes an operation in an existing chaincode. See also chaincode.

trigger

A mechanism that initiates actions. Triggers can be explicitly fired by a user or fired on behalf of a user by an external event source.

true negative

An answer or annotation that is actually incorrect and is predicted to be incorrect.

true positive

An answer or annotation that is actually correct and is predicted to be correct.

trusted profile

A grouping of federated users, compute resources, or both, that are granted the same IAM access. When an identity applies a trusted profile, temporary security credentials are provided for the duration of a session. All identities that are allowed to apply a single profile inherit the same access.

trusted root

A certificate signed by a trusted certificate authority (CA). See also intermediate certificate, certificate authority.

type system

The set of objects that may be discovered by a text analysis engine in a document. The type system defines all possible feature structures in terms of types and features. Any number of different types can be defined in a type system. A type system is domain and application specific.

U

Uniform Resource Identifier (URI)

A unique address that is used to identify content on the web. The most common form of URI is the web page address, which is a particular form or subset of URI called a Uniform Resource Locator (URL). A URI typically describes how to access the resource, the computer that contains the resource, and the location of the resource on that computer.

Uniform Resource Locator (URL)

The unique address of an information resource that is accessible in a network such as the Internet. The URL includes the abbreviated name of the protocol used to access the information resource and the information used by the protocol to locate the information resource. See also host, subdomain, custom domain, domain, route.

unit of measure (UOM)

A standardized unit that is used to measure a specified property of something.

UOM

See unit of measure.

URI

See Uniform Resource Identifier.

URL

See Uniform Resource Locator.

user

- An IBMid or SoftLayer ID that is used as a person's identity in an account.
- A participant in a blockchain network that has indirect access to the ledger through a trust relationship to an existing member.

user registry

A collection of user information, such as user IDs and passwords, that is used as the basis for security control by a system such as a web application server.

V

virtual

Pertaining to not physically existing as such but made by software to appear to do so.

virtual local area network (VLAN)

A logical association of switch ports based upon a set of rules or criteria, such as Medium Access Control (MAC) addresses, protocols, network address, or multicast address. This concept permits the LAN to be segmented again without requiring physical rearrangement.

virtual machine (VM)

A software implementation of a machine that executes programs like a real machine. See also virtual server.

virtual private cloud (VPC)

A virtual network that is tied to a private user account and isolated from other networks in a public cloud. Only authorized users can access virtual private cloud resources, which include virtual servers, storage, and subnets.

virtual private network (VPN)

A private connection between two endpoints, even when the data is transferred across a public network. Usually, a VPN is used in combination with security methods, such as authentication and encryption, to provide maximum data security and privacy.

virtual server

A server that shares its resources with other servers to support applications. See also virtual machine.

VLAN

See virtual local area network.

VM

See virtual machine.

volume

A fixed amount of physical or virtual storage on a data storage medium.

VPC

See virtual private cloud.

VPN

See virtual private network.

VPN as a service

A private connection between two endpoints, which remains private and can be encrypted even when the data is transferred across a public network.

W

WAR

See web archive.

WAR file

See web archive.

web app

See web application.

web application (web app)

An application that is accessible by a web browser and that provides some function beyond static display of information, for instance by allowing the user

to query a database. Common components of a web application include HTML pages, JSP pages, and servlets. See also app.

web archive (WAR)

A compressed file format, defined by the Java EE standard, for storing all the resources required to install and run a web application in a single file.

workload

A set of applications, services, or capabilities that consumes compute, network, storage, or other resources to complete a business objective.

workspace

A context that contains a collection of artifacts that a user with appropriate permission can modify.

Z

zone

A location within a region that acts as an independent fault domain and has decreased latency to other zones in the region. See also multizone region, single-campus multizone region.

Accessibility features for IBM Cloud

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Accessibility features

IBM Cloud® includes the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

IBM Cloud uses the W3C Standard, WAI-ARIA 1.1, to ensure compliance to ensure compliance to US Section 508, Web Content Accessibility Guidelines (WCAG) 2.1, and EN 301 549. To take advantage of accessibility features, use the latest release of your screen reader in combination with the latest Firefox web browser that is supported by this product.

The IBM Cloud online product documentation and the IBM Cloud user interface framework is enabled for accessibility.

Interface information

Review the following information about the IBM Cloud user interface:

- If you are using a screen reader with the IBM Cloud web user interface or product documentation, use the latest version of Firefox with the latest release of the screen reader.
- If you are using keyboard operation only, ensure that your MacOS setting is enabled for `Keyboard navigation: Use keyboard navigation to move focus between controls`.
- The IBM Cloud user interfaces do not have content that flashes 2 - 55 times per second.
- The IBM Cloud web user interfaces rely on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use a user's system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.
- The IBM Cloud web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Related accessibility information

The IBM Cloud web user interface accessibility compliance status is specifically for the IBM Cloud product platform. There are subsections of the user interface that are owned by third-party products or services that host content within the platform, for which the IBM Cloud compliance record does not maintain or own the accessibility compliance status. If you are accessing any user interface or documentation for a service, you must request the compliance statements for that service. For example, if you are using an interface for IBM Cloud Kubernetes Service, the administration console for a local or dedicated environment, or an IoT service, you must request product accessibility information for that interface or documentation.

The IBM Cloud documentation accessibility compliance is specifically for the IBM Cloud core platform information, and it does not extend to any services. The available documentation for IBM Cloud is managed and reported in the IBM Cloud product accessibility information that is available upon request. For compliance status for any service, you must request product accessibility information.

In addition to standard IBM help desk and support websites, IBM has established a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service 800-IBM-3383 (800-426-3383) (within North America)

IBM and accessibility

For more information about the commitment that IBM has to accessibility, see IBM Accessibility.

Legal resources

IBM Cloud Terms of Use

Agreement Terms

The following terms apply to an IBM Cloud® account and use of the IBM Cloud Services:

1. The IBM Cloud Service Description available at https://www.ibm.com/support/customer/csol/terms/?id=i126-6605.
2. Additional Service Descriptions (SDs) for each Cloud Service ordered using your IBM Cloud account, available in the IBM Cloud catalog or on the IBM Terms website. Additional SDs contain links to the applicable Data Sheet for each Cloud Service.
3. One of the following base agreements based upon selected account type and country.
    - IBM Cloud International B.V. (formerly SoftLayer Dutch Holding B.V.) Cloud Services Agreement available at https://www.ibm.com/support/customer/csol/terms/?id=Z126-6304_IBMCIBV which applies to:
        - Enterprise Savings Plan accounts for Clients with business addresses in France, Italy, Spain, Brazil, India, Japan, Israel, or Mexico, invoiced by IBM Cloud International B.V.
        - U.S. Dollar Pay as you Go (PayGo) credit card accounts in all countries **except** Australia, Austria, Belgium, Canada, Denmark, Finland, Germany, Ireland, Luxembourg, Netherlands, New Zealand, Norway, Sweden, Switzerland, United Kingdom, or the United States of America.
        - Free/Lite Accounts.
    - The country-specific IBM Cloud Service Agreement found at https://www.ibm.com/support/customer/csol/terms/?id=Z126-6304 applicable to Client's business address which applies to all other accounts contracted with IBM local country entities, including:
        - Pay as you go (PayGo) credit card accounts for Australia, Austria, Belgium, Canada, Denmark, Finland, Germany, Ireland, Luxembourg, Netherlands, New Zealand, Norway, Sweden, Switzerland, United Kingdom, or the United States of America.
        - Invoiced (non-credit card) and local country currency accounts of any type, except for the Enterprise Savings Plan accounts in countries listed above.
    - The IBM Passport Advantage Agreement you accepted upon enrollment in the Passport Advantage Program available at https://www.ibm.com/support/customer/csol/terms/?id=i126-5948.
    - A business partner agreement covering the use of Cloud Services (e.g. use of an IBM Cloud account for an embedded service or downstream distribution, subject to the terms of the IBM Embedded Solution Agreement, as approved by IBM).
    - Another agreement with equivalent cloud services terms signed between the parties.

Exchange Rate Policy

IBM Cloud adjusts the exchange rates used to provide pricing in non-US dollar currencies on a monthly basis. Charges for cloud services are based on US dollars. Non-US dollar pricing is calculated by converting the US dollar rate by using market exchange rates published by leading financial institutions. Market exchange rates will be adjusted monthly, except as prohibited or controlled by applicable law.

IBM Business Associate Addendum

If you or your company is a covered entity as defined by the US Health Insurance Portability and Accountability Act (HIPAA) and intend to order Cloud Services that might process protected health information (PHI), you must accept the IBM® Business Associate Addendum (BAA) available at https://www.ibm.com/support/customer/csol/terms/?cat=baa. The BAA can be digitally accepted as described in Enabling the HIPAA Supported setting.

EU Data Act

If you are a client based in the European Union and are planning to migrate your data out of IBM Cloud, you are entitled to reduced data egress charges under the terms of the EU Data Act (Regulation (EU) 2023/2854). Please contact your IBM or IBM Business Partner representative, or create a new case in the IBM Cloud Support Center for further information and assistance.

Abuse reporting and legal contact information

For subpoenas or for reporting abuse on IBM Cloud, contact the following:

```
IBM Cloud c/o SoftLayer Inc*.
14001 North Dallas Parkway, Suite M100
Dallas, TX 75240
214.442.0600 Main
214.442.0605 Abuse
214.442.0612 Fax
Subpoenas subpoenas@softlayer.com
```

```
Abuse abuse@softlayer.com
```

* SoftLayer Inc. is an IBM Company

Australian Online Safety Act

For residents of Australia, be aware of your additional safety rights:

- https://www.esafety.gov.au/report
- https://onlinesafety.org.au/codes/

IBM Cloud Notices

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

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

IBM may have patents or pending patent applications covering subject matter that is described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing IBM Corporation North Castle Drive, MD-NC119 Armonk, NY 10504-1785 US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing Legal and Intellectual Property Law IBM Japan Ltd. 19-21, Nihonbashi-Hakozakicho, Chuo-ku Tokyo 103-8510, Japan

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

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

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

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

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information, which has been exchanged, should contact:

IBM Director of Licensing IBM Corporation North Castle Drive, MD-NC119 Armonk, NY 10504-1785 US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

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

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

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

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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

COPYRIGHT LICENSE:

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

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. 2019.

SANCTIONS & EMBARGOES

The United States prohibits most commercial transactions and other dealings with a handful of countries for foreign policy and national security reasons. These prohibitions – formally known as trade and economic sanctions or comprehensive embargoes – currently target Iran, Cuba, North Korea, and Syria. IBM Cloud implements network-wide blocking of communications to IP addresses which originate from countries subject to those trade and economic sanctions. This policy affects all IBM Cloud data center locations and environments, and blocks access by IP addresses registered to countries subject to U.S. trade and economic sanctions.

To identify IP addresses associated with U.S. embargoed countries, IBM Cloud references the Geo IP2 Database maintained by MaxMind). IBM Cloud updates its IP address country mapping on a quarterly basis.

Note: This restriction does not apply to content that is not prohibited under U.S. sanctions and embargoes. Please email IBM_Cloud_Embargo_Exemption_Request@wwpdl.vnet.ibm.com if you believe your content is not prohibited under U.S. trade and economic sanctions.

Notices and information for excluded components

IBM Cloud console

This section consists of the following:

- CC-BY-2.5
- CC-BY-3.0
- CC-BY-4.0
- CC-BY-SA-3.0
- CC-BY-SA-4.0

Creative Commons Attribution 2.5 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by/2.5/legalcode.

JCIP-ANNOTATIONS [HttpCore] (Copyright not found) (https://downloads.apache.org/httpcomponents/httpcore/source/)

Creative Commons Attribution 3.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by/3.0/legalcode.

(FONTAWESOME - DOCS) MONGO-R3.4.1/SRC/THIRD_PARTY/WIREDTIGER/TOOLS/WTSTATS/TEMPLATE/APP/LESS/FONTAWESOME.LESS [MongoDB] (Copyright not found)(http://registry.npmjs.org/mongodb/-/mongodb-3.4.1.tgz)

SPDX [npm] (Copyright not found)(https://github.com/nodejs/node/tree/v12.13.0/deps/npm)

Creative Commons Attribution 4.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by/4.0/legalcode.

CANIUSE-LITE (Copyright (c) Alexis Deveria)(http://registry.npmjs.org/caniuse-lite/-/caniuse-lite-1.0.30000974.tgz)

Creative Commons Attribution Share Alike 3.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by-sa/3.0/legalcode.

T/MAXMIND-DB/* [libmaxminddb-dev] (Copyright: 2013-2014 MaxMind, Inc.)(https://packages.debian.org/stretch/libmaxminddb-dev)

KIMBIE-HIGHLIGHT.JS [highlight.js] (Copyright not found)(http://registry.npmjs.org/highlight.js/-/highlight.js-9.17.1.tgz)

Creative Commons Attribution Share Alike 4.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by-sa/4.0/legalcode.

GLOB LOGO [Glob] (No copyright found)(http://registry.npmjs.org/glob/-/glob-7.1.6.tgz)

LOGO [Glob] (No copyright found)(http://registry.npmjs.org/glob/-/glob-7.1.4.tgz)

BSS

This section consists of the following:

- CC-BY-3.0
- CC-BY-SA-2.5
- CC-BY-SA-3.0

Creative Commons Attribution 3.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by/3.0/legalcode.

SPDX STANDARD. [spdx-expression-parse] (Copyright not found)(http://registry.npmjs.org/spdx-expression-parse/-/spdx-expression-parse-1.0.4.tgz)

BOOTSRAP DOCS [Bootstrap] (Copyright 2011-2018 Twitter, Inc.)(http://registry.npmjs.org/bootstrap/-/bootstrap-3.4.0.tgz)

Creative Commons Attribution Share Alike 2.5 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by-sa/2.5/legalcode.

TEMPLATE:CSSDATA [postcss-reduce-initial] (Copyright not found)(http://registry.npmjs.org/postcss-reduce-initial/-/postcss-reduce-initial-4.0.3.tgz)

Creative Commons Attribution Share Alike 3.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by-sa/3.0/legalcode.

CODE_OF_CONDUCT.MD [enzyme-to-json] (Copyright not found)(http://registry.npmjs.org/enzyme-to-json/-/enzyme-to-json-3.4.3.tgz)

IBM Cloud IAM-Token Service

This section consists of the following:

- CC-BY-3.0

Creative Commons Attribution 3.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by/3.0/legalcode.

MEMORY_SIZE_SYSTEM [Redis] (No Copyright found)(https://github.com/antirez/redis/tree/5.0.0)

ZMALLOC.C [Redis] (No Copyright found) (http://download.redis.io/releases/redis-5.0.6.tar.gz)

IBM Cloud Developer Console

View the Notices file for the IBM Cloud Developer Console.

This section consists of the following:

- AGPL v3

- CC-BY-2.5
- CC-BY-3.0
- CC-BY-4.0
- CC-BY-SA-1.0
- Creative Commons Attribution Share-Alike 2.0
- Creative Commons Attribution Share-Alike 2.5
- Creative Commons Attribution Share-Alike 4.0
- General Public License V2
- General Public License V3
- GNU FREE DOCUMENTATION LICENSE V1.3
- RHEL Universal Base Image
- SIL OFL 1.1
- Unique

GNU AFFERO GENERAL PUBLIC LICENSE, VERSION 3

The Program includes some or all of the following licensed to you as Separately Licensed Code under the GNU Affero General Public License. For copies of the source code for this software, send an email to devlic@us.ibm.com identifying the IBM product and the AGPL-licensed program for which you are requesting the source code.

GOLANGCI-LINT

Creative Commons Attribution 2.5 license

The Program includes some or all of the following that IBM obtained under the Creative Commons Attribute 2.5 License:

GuardedBy.java (https://github.com/stephenc/jcip-annotations/blob/master/src/main/java/net/jcip/annotations/GuardedBy.java), JCIP-ANNOTATIONS (https://github.com/stephenc/jcip-annotations), net.jcip.annotations (https://mvnrepository.com/artifact/net.jcip/jcip-annotations)

The URL to the license is https://creativecommons.org/licenses/by/2.5/

Creative Commons Attribution 3.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by/3.0/legalcode

spdx-exceptions (https://github.com/jslicense/spdx-exceptions.json/archive/v2.2.0.zip), SPDX (https://github.com/spdx), SPDX STANDARD [spdx-expression-parse] (https://github.com/spdx/spdx-spec), GO GOPHER gonum/gonum, DOCUMENTATION opencontainers/runc, H2DEMO/TMPL.GO [x/net], POLY google-api-go-client, DIR.TMPL [x/tools], THE IVY MASCOT x/mobile

Creative Commons Attribution 4.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by/4.0/legalcode

heartbeat_test.c (https://packages.debian.org/stretch/libssl1.0.2), NPM-DISPUTES.HTML [npm] (https://github.com/nodejs/node/tree/v12.13.0/deps/npm), README.MD FILE AND DOCS docker/go-metrics, ZSCHEMA-BROWSER-TEST.JS z-schema, README AND DOCS FOLDER containerd/containerd

Creative Commons Attribution Share Alike 1.0 license

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by-sa/1.0/legalcode

NORMALIZE_TEST.GO x/text

Creative Commons Attribution Share Alike 2.0 license

The Program includes some or all of the following that IBM obtained under the Creative Commons Attribute Share-Alike 2.0 License. The URL to the license is https://creativecommons.org/licenses/by-sa/2.0/legalcode

go-digest docs (https://github.com/opencontainers/go-digest/tree/279bed98673dd5bef374d3b6e4b09e2af76183bf), mapstruct documentation (https://github.com/mapstruct/mapstruct/tree/1.2.0.Final), NORMALIZE_TEST.GO x/text

Creative Commons Attribution Share Alike 2.5 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by-sa/2.5/legalcode

NPM-DISPUTES.HTML (https://github.com/nodejs/node/tree/v12.13.0/deps/npm), DISPUTES.MD npm, DISPUTES/INDEX.HTML npm, DISPUTES.7 npm, x/text

Creative Commons Attribution Share Alike 3.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by-sa/3.0/legalcode

NORMALIZE_TEST.GO x/text

Creative Commons Attribution Share Alike 4.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by-sa/4.0/legalcode

GO-METRICS (DOC) [moby/moby] (Not Required), DOCS docker/spdystream

General Public License V2

The Program includes some or all of the following licensed to you as Separately Licensed Code under the GNU General Public License. For copies of the source code for this software, send an email to devlic@us.ibm.com identifying the IBM product and the GPL-licensed program for which you are requesting the source code.

default-stylesheet.js, Dockerfile.xml, scanelf, tap-driver.sh, OPENJDK,DEBIAN-ARCHIVE-KEYRING ,DOCKER.INITD,SELINUX, UPX [OpenJDK], GOLANGCI/CHECK, OPENRC [moby/moby], ERRNO-BASE.H [modernc.org/cc], TUNED-PROFILES-ORIGIN-NODE.7 [openshift/origin]

GNU General Public License, Version 3

The Program includes some or all of the following licensed to you as Separately Licensed Code under the GNU General Public License. For copies of the source code for this software, send an email to devlic@us.ibm.com identifying the IBM product and the GPL-licensed program for which you are requesting the source code

OPENPEEDEEP/DEPGUARD, LIMITS.H [modernc.org/cc]

GNU Free Documentation License

The Program includes some or all of the following available under the GNU Free Documentation License v1.3. For copies of the source code for this software, send an email to devlic@us.ibm.com identifying the IBM product and the GPL-licensed program for which you are requesting the source code.

calc (GFDL v1.3) [modernc.org/golex]

RHEL Universal Base Image (UBI) v7

RedHat UBI v7 has been taken under EULA > https://www.redhat.com/licenses/eulas

The Program maybe provided with third party software programs subject to their own license terms. The license terms either accompany the third party software programs or, in some instances, maybe viewed at https://access.redhat.com/containers/.

SIL OFL 1.1

The Program includes some or all of the following that IBM obtained under the SIL Open Font License version 1.1:

Font Awesome 4.1.0 (by @davegandy), Orbitron Font (Copyright (c) 2009, Matt McInerney (matt@pixelspread.com)), Oswald Font (Copyright (c) 2011-2012, Vernon Adams (vern@newtypography.co.uk)) ORBITRON FONTS [smartystreets/goconvey], FONT [openshift/origin], FONT AWESOME FONTS [smartystreets/goconvey]

SIL Open Font License 1.1

PREAMBLE The goals of the Open Font License (OFL) are to stimulate worldwide development of collaborative font projects, to support the font creation efforts of academic and linguistic communities, and to provide a free and open framework in which fonts may be shared and improved in partnership with others.

The OFL allows the licensed fonts to be used, studied, modified and redistributed freely as long as they are not sold by themselves. The fonts, including any derivative works, can be bundled, embedded, redistributed and/or sold with any software provided that any reserved names are not used by derivative works. The fonts and derivatives, however, cannot be released under any other type of license. The requirement for fonts to remain under this license does not apply to any document created using the fonts or their derivatives.

DEFINITIONS "Font Software" refers to the set of files released by the Copyright Holder(s) under this license and clearly marked as such. This may include source files, build scripts and documentation.

"Reserved Font Name" refers to any names specified as such after the copyright statement(s).

"Original Version" refers to the collection of Font Software components as distributed by the Copyright Holder(s).

"Modified Version" refers to any derivative made by adding to, deleting, or substituting -- in part or in whole -- any of the components of the Original Version, by changing formats or by porting the Font Software to a new environment.

"Author" refers to any designer, engineer, programmer, technical writer or other person who contributed to the Font Software.

PERMISSION & CONDITIONS Permission is hereby granted, free of charge, to any person obtaining a copy of the Font Software, to use, study, copy, merge, embed, modify, redistribute, and sell modified and unmodified copies of the Font Software, subject to the following conditions:

1. Neither the Font Software nor any of its individual components, in Original or Modified Versions, may be sold by itself.
2. Original or Modified Versions of the Font Software may be bundled, redistributed and/or sold with any software, provided that each copy contains the above copyright notice and this license. These can be included either as stand-alone text files, human-readable headers or in the appropriate machine-readable metadata fields within text or binary files as long as those fields can be easily viewed by the user.
3. No Modified Version of the Font Software may use the Reserved Font Name(s) unless explicit written permission is granted by the corresponding Copyright Holder. This restriction only applies to the primary font name as presented to the users.
4. The name(s) of the Copyright Holder(s) or the Author(s) of the Font Software shall not be used to promote, endorse or advertise any Modified Version, except to acknowledge the contribution(s) of the Copyright Holder(s) and the Author(s) or with their explicit written permission.
5. The Font Software, modified or unmodified, in part or in whole, must be distributed entirely under this license, and must not be distributed under any other license. The requirement for fonts to remain under this license does not apply to any document created using the Font Software.

TERMINATION This license becomes null and void if any of the above conditions are not met.

DISCLAIMER THE FONT SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF COPYRIGHT, PATENT, TRADEMARK, OR OTHER RIGHT. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, INCLUDING ANY GENERAL, SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF THE USE OR INABILITY TO USE THE FONT SOFTWARE OR FROM OTHER DEALINGS IN THE FONT SOFTWARE.

UNIQUE (SLC REQD)

Debian, Alpine OS & Its Components

IBM Cloud CLI

View the Notices file for the IBM Cloud CLI.

This section consists of the following:

- CC-BY-2.5
- CC-BY-3.0
- CC-BY-4.0
- CC-BY-SA-1.0
- CC-BY-SA-2.0
- CC-BY-SA-2.5
- CC-BY-SA-3.0
- CC-BY-SA-4.0
- GPL V2
- GNU FREE DOCUMENTATION LICENSE
- REDHAT UNIVERSAL BASE IMAGE
- SIL OPEN FONT LICENSE 1.1
- UNIQUE (SLC REQD)

Creative Commons Attribution 2.5 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by/2.5/legalcode.

- GUARDEDBY.JAVA JSR305

- JCIP-ANNOTATIONS HttpCore
- IMMUTABLE.JAVA JSR305
- NOTTHREADSAFE.JAVA JSR305

Creative Commons Attribution 3.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by/3.0/legalcode

- BOOTSTRAP DOCS
- SITE-POLICIES golang/net, DIR.TMPL - SITE POLICIES golang/tools
- CONTENT-SITE-POLICIES#RESTRICTIONS golang.org/x/tools, DOCS docker/libtrust, GODOCS btoa (Not Required), GOPROXY-JQUERY-VERSION elazarl/goproxy(Not Required)
- WEBPAGECONTENTS-HTML/TESTDATA/GO1.HTML github.com/hashicorp/go.net (Not Required), GO1.HTML github.com/hashicorp/go.net (Not Required), GODOC.HTML (no copyright), PHP_MAN.HTML elazarl/goproxy (Not Required), GOLANG-ORG.YAML github.com/dnaeon/go-vcr (Not Required), POLY-GEN.GO (no copyright), FONT AWESOME ANGularJS (Not Required), ABOUTIVY.HTML (no copyright), DIR.TMPL (no copyright), SPRITES (no copyright), GOPHER gonum.org/v1/gonum (Not Required), GOPHER IMAGE BY RENEE FRENCH Go (Copyright (c) Renee French)
- PAGE CONTENT TMPL.GO golang.org/x/net (Copyright 2017 The Go Authors)
- FOOTER (Copyright © OpenJS Foundation.) `(codedrop-06072022/source/resource-catalog-cli.zip/resource-catalog-cli-d471ddd3d38302c9c93e27d8c62b242563d36c1c/vendor/github.com/fatih/color/Gopkg.lock)`, ICONS github.com/smartystreets/goconvey (Not Required)

Creative Commons Attribution 4.0 Non-IBM-License-Information

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by/4.0/legalcode.

- SPDYSTREAM/DOCS docker/spdystream, FIXTURES [go-openapi/spec] (Not Required)
- DOCS containerd/containerd (Not Required), EXPANDER_TEST.GO go-openapi/spec (Not Required)

Creative Commons Attribution Share Alike 1.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by-sa/1.0/legalcode.

TXT_RU golang/text

Creative Commons Attribution Share Alike 2.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by-sa/2.0/legalcode

- MASTER SHAKE PHOTO BY SOL GRUNDY (MASTER_SHAKE.JPG) rest-client (Ruby Gem)
- TXT_KR golang/text

Creative Commons Attribution Share Alike 2.5 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by-sa/2.5/legalcode

TXT_CN golang/text

Creative Commons Attribution Share Alike 3.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by-sa/3.0/legalcode

- envoyproxy/protoc-gen-validate
- TXT_VN golang/text
- IMAGES jQuery-ui (Not Required)

Creative Commons Attribution Share Alike 4.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by-sa/4.0/legalcode

- README.MD docker-ce-cli
- GO-METRICS DOCS docker-ce-cli, DOCS go-digest, GLOB LOGO Glob

GNU General Public License, Version 2

The Program includes some or all of the following licensed to you as Separately Licensed Code under the GNU General Public License. For copies of the source code for this software, send an email to devlic@us.ibm.com identifying the IBM product and the GPL-licensed program for which you are requesting the source code.

- DEBIAN-ARCHIVE-KEYRING
- DEFAULTER.JS [Adm-zip]
- ASM-GENERIC [modernc.org/cc], /SBIN/OPENRC-RUN [docker/engine], GHOSTSCRIPT [modernc.org/cc]

GNU General Public License, Version 3

The Program includes some or all of the following licensed to you as Separately Licensed Code under the GNU General Public License. For copies of the source code for this software, send an email to devlic@us.ibm.com identifying the IBM product and the GPL-licensed program for which you are requesting the source code.

V2/HEADERS [modernc.org/cc], TESTDATA/GCC [modernc.org/cc]

GNU Free Documentation License

The Program includes some or all of the following available under the GNU Free Documentation License v1.3. For copies of the source code for this software, send an email to devlic@us.ibm.com identifying the IBM product and the GPL-licensed program for which you are requesting the source code.

hackerslab.org (tla), tla-1.3.4/src/libneon/doc (tla), documentation (json-glib), Package documentation (Expat), doc (autogen), libtasn1.info

RedHat Universal Base Image v8.4 (OOTPA)

RedHat UBI v7, v7.6 and 8.0 has been taken under EULA > https://www.redhat.com/licenses/eulas.

The Program may be provided with third party software programs subject to their own license terms. The license terms either accompany the third party software programs or, in some instances, maybe may be viewed at https://access.redhat.com/containers/.

SIL OPEN FONT LICENSE 1.1

The Program includes some or all of the following works licensed under the SIL OPEN FONT LICENSE 1.1.

ORBITRON/OSWALD FONT [smartystreets/goconvey], FONT AWESOME FONT [smartystreets/goconvey]

PREAMBLE

The goals of the Open Font License (OFL) are to stimulate worldwide development of collaborative font projects, to support the font creation efforts of academic and linguistic communities, and to provide a free and open framework in which fonts may be shared and improved in partnership with others.

The OFL allows the licensed fonts to be used, studied, modified and redistributed freely as long as they are not sold by themselves. The fonts, including any derivative works, can be bundled, embedded, redistributed and/or sold with any software provided that any reserved names are not used by derivative works. The fonts and derivatives, however, cannot be released under any other type of license. The requirement for fonts to remain under this license does not apply to any document created using the fonts or their derivatives.

DEFINITIONS

"Font Software" refers to the set of files released by the Copyright Holder(s) under this license and clearly marked as such. This may include source files, build scripts and documentation.

"Reserved Font Name" refers to any names specified as such after the copyright statement(s).

"Original Version" refers to the collection of Font Software components as distributed by the Copyright Holder(s).

"Modified Version" refers to any derivative made by adding to, deleting, or substituting - in part or in whole - any of the components of the Original Version, by changing formats or by porting the Font Software to a new environment.

"Author" refers to any designer, engineer, programmer, technical writer or other person who contributed to the Font Software.

PERMISSION & CONDITIONS

Permission is hereby granted, free of charge, to any person obtaining a copy of the Font Software, to use, study, copy, merge, embed, modify, redistribute, and sell modified and unmodified copies of the Font Software, subject to the following conditions:

1. Neither the Font Software nor any of its individual components, in Original or Modified Versions, may be sold by itself.
2. Original or Modified Versions of the Font Software may be bundled, redistributed and/or sold with any software, provided that each copy contains the above copyright notice and this license. These can be included either as stand-alone text files, human-readable headers or in the appropriate machine-readable metadata fields within text or binary files as long as those fields can be easily viewed by the user.
3. No Modified Version of the Font Software may use the Reserved Font Name(s) unless explicit written permission is granted by the corresponding Copyright Holder. This restriction only applies to the primary font name as presented to the users.
4. The name(s) of the Copyright Holder(s) or the Author(s) of the Font Software shall not be used to promote, endorse or advertise any Modified Version, except to acknowledge the contribution(s) of the Copyright Holder(s) and the Author(s) or with their explicit written permission.
5. The Font Software, modified or unmodified, in part or in whole, must be distributed entirely under this license, and must not be distributed under any other license. The requirement for fonts to remain under this license does not apply to any document created using the Font Software.

TERMINATION

This license becomes null and void if any of the above conditions are not met.

DISCLAIMER

THE FONT SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF COPYRIGHT, PATENT, TRADEMARK, OR OTHER RIGHT. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, INCLUDING ANY GENERAL, SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF THE USE OR INABILITY TO USE THE FONT SOFTWARE OR FROM OTHER DEALINGS IN THE FONT SOFTWARE.

Unique (SLC REQD)

Alpine Linux (main) Debian (matching Ubuntu main) UBUNTU (MAIN)

IBM Cloud Shell

This section consists of the following:

- CC-BY-1.0
- CC-BY-2.5
- CC-BY-3.0
- CC-BY-4.0
- CC-BY-SA-1.0
- CC-BY-SA-2.0
- CC-BY-SA-2.5
- CC-BY-SA-3.0
- CC-BY-SA-4.0

Creative Commons Attribution 1.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by/1.0/legalcode

SIPHASH NodeJS

Creative Commons Attribution 2.5 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by/2.5/legalcode

JCIP-ANNOTATIONS HttpCore, JAVAX.ANNOTATION.CONCURRENT libjsr305-java

Creative Commons Attribution 3.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by/3.0/legalcode

DOCS. jq, FONT AWESOME DOCUMENTATION goconvey, SOUNDMANAGER.TS xterm, H2DEMO golang/net, SPDX-EXCEPTIONS, SPDX STANDARD spdx-expression-parse, FIREWORKS.JPEG Snappy, GO1.HTML hashicorp/go.net,GODOC.HTML [golang.org/x/tools] (Not Required), GOLANG-ORG.YAML dnaeon/go-vcr, GO Go, HTTP2/H2DEMO/TMPL.GO dns, POLY google.golang.org/api, GOPHER IMAGE Go, DIR.TMPL [golang.org/x/tools] (Not Required),

SPRITES golang.org/x/mobile, GOPHER gonum, RUNC-DOCUMENTATION opencontainers/runc

Creative Commons Attribution 4.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by/4.0/legalcode

DOM Typescript, GO-DIGEST-README.MD FILE AND CONTRIBUTING.MD distribution, DOC docker-ce-cli, TEST_CONTAINERD_DOCS Microsoft/hcsshim, CASCADINGSCHEMES.YAML go-openapi/loads, CONTAINERD DOCUMENTATION docker/docker, DATA.GO go-openapi/runtime, DOCS/LOGO.SVG Semver, SWAGGER PETSTORE [go-openapi/validate], EXPANDER_TEST.GO go-openapi/spec, PAPER-100K.PDF Snappy, README AND DOCS FOLDER containerd/containerd

Creative Commons Attribution Share Alike 1.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by-sa/1.0/legalcode

MAP_TEST.GO [x/text], RUSSIAN TEXT golang.org/x/text

Creative Commons Attribution Share Alike 2.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by-sa/2.0/legalcode

JAPANESE TEXT golang.org/x/text, THREEBYTEUTF8 [x/text]

Creative Commons Attribution Share Alike 2.5 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by-sa/2.5/legalcode

CANTONESE TEXT golang.org/x/text, MAP_TEST.GO [x/text]

Creative Commons Attribution Share Alike 3.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by-sa/3.0/legalcode

REGULAR_EXPRESSIONS posix-character-classes, TEXT IN FOREIGN LANGUAGE golang/text, NEP-0012-MISSING-DATA.RST NumPy, VIETNAMESE TEXT golang.org/x/text, RAMDA LOGO RAMda, NEP-0012-MISSING-DATA.RST NumPy, DATA LICENSE (GEOLITE COUNTRY AND CITY DATABASES) geoip, NEP12 [NumPy], MAP_TEST.GO [x/text]

Creative Commons Attribution Share Alike 4.0 Generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by-sa/4.0/legalcode

DOCS kubernetes-sigs/cri-tools, GLOB LOGO Glob, SCREENSHOTS shortid, GO-METRICS DOCS docker/go-metrics, DISPUTES.7 [NodeJS] (Not Required), GLOB NodeJS, GO-DIGEST DOCS Microsoft/hcsshim, GO-DIGEST DOCUMENTATION go-digest, DISPUTES npm, DISPUTES DOC [NodeJS] (Not Required), SCREENSHOTS shortid

IBM Cloud Security and Compliance Center

View the Notices file for the Security and Compliance Center.

This section consists of the following:

- CC-BY-3.0
- CC-BY-SA-2.0
- GPL V2
- GPL V3
- SIL OPEN FONT LICENSE 1.1

Creative Commons Attribution 2.5 generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by/2.5/legalcode.

JAVAX.ANNOTATION.CONCURRENT [JSR305] ((c) 2016 SPDX Workgroup a Linux Foundation Collaborative Project) (https://repo1.maven.org/maven2/com/google/code/findbugs/jsr305/1.3.9/)

Creative Commons Attribution 3.0

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by/3.0/legalcode.

- PHP_MAN.HTML [elazarl/goproxy (go)] (Copyright (c) 2012 Elazar Leibovich) (https://github.com/elazarl/goproxy/tree/49ad98f6dac1e021cff612afa5a9d95843efde52)
- SPDX [npm] (Copyright (c) npm, Inc. and Contributors)(http://registry.npmjs.org/npm/-/npm-6.14.6.tgz)
- SITE-POLICIES [golang/tools (go)] (Copyright (c) 2009 The Go Authors) (https://github.com/golang/tools/archive/4735f53fb37db012e63d1a62914417980fb3c023.zip)

Creative Commons Attribution Share Alike 2.0 generic

The Program includes some or all of the following works licensed under the Creative Commons. The URL to the license is https://creativecommons.org/licenses/by-sa/2.0/legalcode.

DISPUTES [npm] (Copyright (c) npm, Inc. and Contributors)(http://registry.npmjs.org/npm/-/npm-6.14.6.tgz)

SIL Open Font License 1.1

The Program includes some or all of the following works licensed under the SIL Open Font License 1.1.

PREAMBLE The goals of the Open Font License (OFL) are to stimulate worldwide development of collaborative font projects, to support the font creation efforts of academic and linguistic communities, and to provide a free and open framework in which fonts may be shared and improved in partnership with others.

The OFL allows the licensed fonts to be used, studied, modified, and redistributed freely as long as they are not sold by themselves. The fonts, including any derivative works, can be bundled, embedded, redistributed and/or sold with any software provided that any reserved names are not used by derivative works. The fonts and derivatives, however, cannot be released under any other type of license. The requirement for fonts to remain under this license does not apply to any document created using the fonts or their derivatives.

DEFINITIONS "Font Software" refers to the set of files released by the Copyright Holder(s) under this license and clearly marked as such. This may include source files, build scripts and documentation.

"Reserved Font Name" refers to any names specified as such after the copyright statement(s).

"Original Version" refers to the collection of Font Software components as distributed by the Copyright Holder(s).

"Modified Version" refers to any derivative made by adding to, deleting, or substituting -- in part or in whole -- any of the components of the Original Version, by changing formats or by porting the Font Software to a new environment.

"Author" refers to any designer, engineer, programmer, technical writer, or other person who contributed to the Font Software.

PERMISSION AND CONDITIONS Permission is hereby granted, free of charge, to any person obtaining a copy of the Font Software, to use, study, copy, merge, embed, modify, redistribute, and sell modified and unmodified copies of the Font Software, subject to the following conditions:

1. Neither the Font Software nor any of its individual components, in Original or Modified Versions, may be sold by itself.
2. Original or Modified Versions of the Font Software may be bundled, redistributed and/or sold with any software, provided that each copy contains the above copyright notice and this license. These can be included either as stand-alone text files, human-readable headers or in the appropriate machine-readable metadata fields within text or binary files as long as those fields can be easily viewed by the user.
3. No Modified Version of the Font Software may use the Reserved Font Name(s) unless explicit written permission is granted by the corresponding Copyright Holder. This restriction only applies to the primary font name as presented to the users.
4. The name(s) of the Copyright Holder(s) or the Author(s) of the Font Software shall not be used to promote, endorse or advertise any Modified Version, except to acknowledge the contribution(s) of the Copyright Holder(s) and the Author(s) or with their explicit written permission.
5. The Font Software, modified or unmodified, in part or in whole, must be distributed entirely under this license, and must not be distributed under any other license. The requirement for fonts to remain under this license does not apply to any document created using the Font Software.

TERMINATION This license becomes null and void if any of the above conditions are not met.

DISCLAIMER THE FONT SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF COPYRIGHT, PATENT, TRADEMARK, OR OTHER RIGHT. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, INCLUDING ANY

GENERAL, SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF THE USE OR INABILITY TO USE THE FONT SOFTWARE OR FROM OTHER DEALINGS IN THE FONT SOFTWARE.

- FONTAWESOME [smartystreets/goconvey] (Copyright (c) 2016 SmartyStreets, LLC)(https://github.com/smartystreets/goconvey/tree/v1.6.4)
- OFL.TXT [smartystreets/goconvey] (Copyright (c) 2016 SmartyStreets, LLC)(https://github.com/smartystreets/goconvey/tree/v1.6.4)

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered 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.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

The Android robot is reproduced or modified from work created and shared by Google and used according to terms described in the Creative Commons 3.0 Attribution License.

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

Node.js is a trademark of Joyent, Inc. and is used with its permission. This documentation is not formally endorsed by or affiliated with Joyent.

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

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

VMware, the VMware logo, VMware Cloud Foundation, VMware Cloud Foundation Service, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

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

Red Hat®, JBoss®, OpenShift®, Fedora®, Hibernate®, Ansible®, CloudForms®, RHCA®, RHCE®, RHCSA®, Ceph®, and Gluster® are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

Salesforce is a trademark of Salesforce.com, Inc.

Box is a trademark of Box, Inc.

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

Terminology

While IBM values the use of inclusive language, terms that are outside of IBM's direct influence are sometimes required for the sake of maintaining user understanding. As other industry leaders join IBM in embracing the use of inclusive language, IBM will continue to update the documentation to reflect those changes.

IP Address Policy

1. IP Address Ownership

International Business Machines Corporation or its subsidiaries that (IBM®) own any IP addresses assigned to Client, which are to be used only with the Cloud Services. Use of the IP addresses terminates upon expiration or termination of a service, at which time the IP addresses will be reclaimed and may be re-issued to other customers in the future.

2. Client-Provided IP Addresses

Clients may provide their own Regional Internet Registry (RIR) issued IP addresses, subject to certain technical limitations and verification of ownership, by submitting a support case and providing a Letter of Authority to specify Client-owned IP prefixes. If IBM approves a request those IP addresses will be announced via Border Gateway Protocol (BGP) on the Client's behalf and routed to identified Client servers. Client retains ownership of those IP addresses. At Client's request, or upon cancellation of the services, IBM will cease to announce and route those IP addresses. While IBM will make reasonable efforts to ensure that BGP prefixes are accepted by all upstream ISPs, we cannot guarantee global reachability for Client-owned IPs.

3. Review / Justification / Efficiency Guidelines

Because IPv4 addresses are a scarce resource, RIRs require ISPs to document that they are efficiently utilizing existing assigned addresses and are planning efficient utilization of any addresses being requested. RIR policies and RFC2050 promote conservation and deter wasteful use or stockpiling of IP

space. IBM is required to abide by these policies when IBM requests additional IP addresses to allocate. Therefore, when requested by IBM, Client provides information necessary to enable IBM to obtain RIR-issued IP addresses to support the Cloud Services.

All Client IP address requests are subjected to review by IBM to ensure efficient utilization and are not guaranteed to be approved. During the review, Client is required to provide details about how each IP address would be utilized and technical justification as to why additional IP addresses are needed. IBM may take steps to validate such information, including network scanning and server configuration inspection.

An initial review may take up to 2 business days, and depending on size and complexity, may require additional time for completion. Accounts with open abuse tickets will have their IP requests held for processing until all abuse issues have been resolved.

Service Level Agreements (SLAs)

IBM Cloud® aims to deliver the highest levels of availability and offers a service level agreement (SLA), which provides credits against service charges should a service fail to meet its stated availability target. For more information, see IBM Cloud Service Level Agreements.

Extended support and product release levels

IBM Cloud® is committed to support a set of key products for extended periods. Depending on a product's release level, various levels of support, committed warranties, or service level agreements (SLAs) are available.

Extended support for key products

The following table lists the products that IBM® will not withdraw support for before the indicated date, unless a functional equivalent is provided.

Name	Continued support through
App ID	2026-12-31
IBM Cloudant	2026-12-31
Cloud Functions	2026-12-31
Databases for PostgreSQL	2027-12-31
IBM Db2 Warehouse SaaS	2026-12-31
Event Streams	2027-12-31
Kubernetes Service	2027-12-31
Object Storage	2026-12-31
Power® Virtual Server	2026-12-31

Product support

For products that are not included in the previous table, IBM will provide at least 30 days notice of its intention to discontinue or replace a generally available product. If no replacement is provided, IBM will use commercially reasonable means to continue the operation and support of previously deployed instances of the product for a period of at least 12 months.

For more information about the terms of use for products, see IBM Cloud Terms.

Product release levels

Products are available in the IBM Cloud catalog at various release levels, including general availability (GA), select availability, beta, or experimental. The catalog also notes deprecated products, which are at the end of their lifecycle.

- GA products are widely available for sale and delivery to customers or channels, usually across multiple geographies.
- Select availability products are production-ready, available for sale, and accessible to select customers.
- Beta products are made solely available for evaluation and testing purposes. There are no warranties, SLAs, or support provided and beta products are not intended for production use.
- Experimental products are made solely available for evaluation and testing purposes, and might be unstable or not compatible with previous

versions. An experimental product can be discontinued with short notice. There are no warranties, SLAs, or support provided, and experimental products are not intended for production use.

- Deprecated products are in the process of being withdrawn from service and are eligible to be removed after the deprecation period.

IBM Cloud price changes

The following page outlines the price changes for IBM Cloud®.

January 2025 price changes

Effective 1 January 2025, IBM Cloud is making price changes for infrastructure as a service (IaaS) and platform as a service (PaaS) services, IBM Cloud Object Storage, R1Soft Backup Service, IBM Cloud Security and Compliance Center Workload Protection cloud security posture management, and IBM Key Protect. See the following sections for additional details.

Premium changes by location for IaaS

Location premiums for most non-US multizone regions (MZRs) will remain unchanged. The premium for Sao Paulo (`br-sao`) will increase by 3 percentage points.

MZR locations	Data centers	Current location premium	New location premium	Effective increase
Dallas (`us-south`)	DAL10, DAL12, DAL13, DAL14	0%	0%	No increase
Washington DC (`us-east`)	WDC04, WDC06, WDC07	0%	0%	No increase
Toronto (`ca-tor`)	TOR01, TOR04, TOR05	6%	6%	No increase
London (`eu-gb`)	LON04, LON05, LON06	13%	13%	No increase
Frankfurt (`eu-de`)	FRA02, FRA04, FRA05	16%	16%	No increase
Madrid (`eu-es`)	MAD02, MAD04, MAD05	16%	16%	No increase
Osaka (`jp-osa`)	OSA21, OSA22, OSA23	20%	20%	No increase
Tokyo (`jp-tok`)	TOK02, TOK04, TOK05	20%	20%	No increase
Sydney (`au-syd`)	SYD01, SYD04, SYD05	20%	20%	No increase
Sao Paulo (`br-sao`)	SAO01, SAO04, SAO05	29%	32%	2.3%

IaaS MZR location premium changes

Location premiums for classic data centers will increase as follows:

Classic data centers	Current location premium	New location premium	Effective increase
DAL08	11%	11%	No increase
DAL09	0%	6%	6.0%
SJC03	0%	6%	6.0%
SJC04	0%	6%	6.0%
WDC03	11%	11%	No increase
AMS03	6%	13%	6.6%

MON01	6%	13%	6.6%
LON02	13%	20%	6.2%
MIL01	16%	22%	5.2%
PAR01	16%	22%	5.2%
CHE01	20%	20%	No increase
SNG01	20%	20%	No increase

Classic data center location premium changes

Note: Existing VMware, Power Virtual Server (including SAP on Power Virtual Server), High Performance Computing, and SAP deployments made before 1 January 2025 will be exempt from these increases and prices will remain unchanged.

Premium changes by location for PaaS services

PaaS services outside of the United States will adopt the same location premiums as IaaS services shown in the previous section.

Excluded from this change are the following services:

- IBM Cloud Kubernetes Service and Red Hat OpenShift on IBM Cloud
- IBM Cloud Logs
- IBM Cloud Activity Tracker Deprecated
- IBM Cloud Code Engine
- IBM Cloud Log Analysis Deprecated
- IBM Cloud Monitoring
- IBM Cloud Databases for Elasticsearch
- IBM Cloud Databases for EnterpriseDB
- IBM Cloud Databases for etcd
- IBM Cloud Databases for MongoDB
- Databases for MySQL
- IBM Cloud Databases for PostgreSQL
- IBM Cloud Databases for Redis
- IBM Cloud Messages for RabbitMQ

Included in this change are the following services:

- IBM Cloud® App Configuration
- IBM Cloud App ID
- IBM Cloudant
- IBM Cloud Container Registry
- IBM Cloud Continuous Delivery
- IBM Cloud Event Notifications
- IBM Cloud Satellite
- IBM Key Protect
- IBM Event Streams for IBM Cloud
- IBM Cloud Secrets Manager
- IBM Cloud Security and Compliance Center
- IBM Cloud Security and Compliance Center Workload Protection

IBM Cloud Object Storage

Object Storage will introduce charges for Aspera high-speed uploads and adopt new, simpler, flat, and consistent charges for high-speed downloads:

Transfer	Current charges	2025 charges - EU/UK/US/Japan	2025 charges - All other regions
High-speed uploads	$0.00	10 GB free/month 10 GB+ $0.04/GB/month	10 GB free/month 10 GB+ $0.08/GB/month
High-speed downloads	0-50 TB $0.0836/GB Next 100 TB $0.0627/GB 150 TB+ $0.0418/GB	10 GB free/month 10 GB+ $0.04/GB/month	10 GB free/month 10 GB+ $0.08/GB/month

New Object Storage high-speed data transfer charges.

Additionally, Object Storage prices will be increased by 2% in the Sao Paulo (`br-sao`) and Toronto (`ca-tor`) regions. In Chennai (`in-che`), the 2% increase will apply to Archive only.

R1Soft Backup Services

Charges for R1Soft Backup Services will increase as follows:

Tier	Current charges	2025 charges
1 Pack	$7.50	$10.00
5 Pack	$36.25	$40.00
10 Pack	$67.50	$75.00
25 Pack	$143.75	$160.00

R1Soft backup services changes

IBM Cloud Security and Compliance Center

Charges for IBM Cloud Security and Compliance Center Workload Protection cloud security posture management will increase from $8.24/compute instance/month to $18.00/compute instance/month.

Volume discounts are available as follows:

Number of compute instances	Current price/month	2025 price/month
1 - 250	$8.24	$18.00
251 - 500	$7.004	$15.30
501 - 1000	$6.18	$13.50
1001 - 2500	$5.768	$12.60
2501 - 5000	$4.944	$10.80
5000+	$4.523	$9.90

IBM Cloud Security and Compliance Center Workload Protection cloud security posture management changes

Note: All other Security and Compliance Center charges remain unchanged.

IBM Key Protect

Key Protect is making two changes to its pricing structure:

1. The current free allocations of 5 key versions will be eliminated. The standard charge of $1.0764 /key version/month will apply to all keys.
2. Charges for the global resiliency feature will be $100.00/duplicate region/month plus $1.0764/month for each duplicated key.

April 2024 price changes

Effective 1 April 2024, IBM Cloud is making the following price changes for cPanel and Red Hat® Enterprise Linux.

Red Hat Enterprise Linux

Effective 1 April 2024, Red Hat is making changes to Red Hat Enterprise Linux (RHEL) prices and introducing a new tier structure for virtual server instance pricing. IBM Cloud is making the following price changes to reflect the Red Hat changes.

RHEL on Bare Metal Servers

For Bare Metal Servers, RHEL prices will change as follows:

Platform	Offering	Current charges	New charges
VPC Bare Metal	RHEL and RHEL for SAP Applications Hourly	$0.120	$0.132
Classic Bare Metal	RHEL Monthly	$90.00	$99.00
Classic Bare Metal	RHEL for SAP Applications Monthly	$90.00	$99.00
Classic Bare Metal	RHEL for SAP with HA and US Monthly	$448.00	$492.80

RHEL on Bare Metal Servers

RHEL on Virtual Server Instances

Red Hat prices for virtual server instances are changing from a 2-tier per server model to a 3-tier per vCPU model as follows:

Old Model - charges based on server size	Size
Small	1-4 vCPUs
Large	5+ vCPUs

Charges based on server size

New model - charges per vCPU	Size
Small	1-8 vCPUs
Mid	9-127 vCPUs
Large	128+ vCPUs

Charges per vCPU

With this new model, prices effective 1 April 2024 will be as follows:

Platform	Offering	Current charges (per server)	New charges (per vCPU)
Classic	VSI RHEL Hourly	Small: $0.06 Large: $0.12	Small: $0.017 Mid: $0.011 Large: not offered
Classic	VSI RHEL Monthly	Small: $45.00 Large: $93.00	Small: $12.584 Mid: $8.829 Large: not offered
VPC	VSI RHEL for SAP Apps Hourly	Small: $0.06 Large: $0.12	Small: $0.017 Mid: $0.011 Large: $0.009

VPC	VSI RHEL for SAP Apps for HA and US Hourly	Small: $0.14 Large: $0.31	Small: $0.023 Mid: $0.016 Large: $0.015
VPC	VSI RHEL Hourly	Small: $0.06 Large: $0.12	Small: $0.017 Mid: $0.011 Large: $0.009
VMware Shared	VSI RHEL Hourly	Small: $0.06 Large: $0.12	Small: $0.017 Mid: $0.011 Large: $0.009
VMware-aaS	VSI RHEL Hourly	Small: $0.06 Large: $0.12	Small: $0.017 Mid: $0.011 Large: $0.009

Comparison of pricing model changes

cPanel

Effective 1 April 2024, IBM Cloud is making the following price changes for cPanel. These prices reflect a vendor price increase that was made in December 2023.

cPanel Option	Current monthly charge	New charge from 1 April 2024
cPanel/WHM with Fantastico and RVskin Admin Cloud up to 5 Accounts	$14.00	$30.25
cPanel/WHM with Fantastico and RVskin Pro Cloud up to 30 Accounts	$19.00	$42.99
cPanel/WHM with Fantastico and RVskin Plus Cloud up to 50 Accounts	$27.00	$57.00
cPanel/WHM with Fantastico and RVskin Premier Cloud up to 100 Accounts	$48.50	$72.55
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Metal up to 100 Accounts	$48.50	$72.55
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 150 Accounts	$63.50	$81.80
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 200 Accounts	$78.50	$100.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 250 Accounts	$93.50	$120.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 300 Accounts	$108.50	$140.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 350 Accounts	$123.50	$160.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 400 Accounts	$138.50	$180.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 450 Accounts	$153.50	$200.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 500 Accounts	$168.50	$220.99

cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 550 Accounts	$183.50	$240.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 600 Accounts	$198.50	$260.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 650 Accounts	$213.50	$280.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 700 Accounts	$228.50	$300.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 750 Accounts	$243.50	$320.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 800 Accounts	$258.50	$340.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 850 Accounts	$273.50	$360.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 900 Accounts	$288.50	$380.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 950 Accounts	$303.50	$400.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 1000 Accounts	$318.50	$420.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 1050 Accounts	$333.50	$440.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 1100 Accounts	$348.50	$460.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 1150 Accounts	$363.50	$480.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 1200 Accounts	$378.50	$500.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 1250 Accounts	$393.50	$520.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 1300 Accounts	$408.50	$540.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 1350 Accounts	$423.50	$560.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 1400 Accounts	$438.50	$580.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 1450 Accounts	$453.50	$600.99
cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 1500 Accounts	$468.50	$620.99

cPanel/WHM with Fantastico/Softaculous and RVskin Premier Fixed Package up to 2000 Accounts	$618.50	$820.99

cPanel pricing changes

January 2024 price changes

Effective 1 January 2024, IBM Cloud is making the following price changes for IaaS and PaaS services.

IaaS data center premium changes

Prices for IBM Cloud infrastructure services vary from location to location, based on a percentage premium from US base prices. Effective 1 January 2024, these location premiums are changing as follows:

Location	Current premium	New premium	Effective increase
Dallas San Jose Washington DC	0%	0%	No Change
Amsterdam Montreal Toronto	3%	6%	+2.9%
London	7%	13%	+5.6%
Frankfurt Madrid Milan Paris	10%	16%	+5.5%
Osaka Singapore Tokyo	13%	20%	+6.2%
Chennai Sydney	20%	20%	No Change
Sao Paulo	20%	29%	+7.5%

IaaS pricing changes

These new premiums will apply to Bare Metal Servers, Virtual Server Instances, File and Block Storage, and Networking infrastructure, for both classic and VPC offerings.

For IBM Cloud® Object Storage, the premiums that are mentioned in the previous table will affect the Object Storage service only in the Sao Paulo (`br-sao`) region. Furthermore, there will be a global increase of 25% in **Accelerated Archive** prices and a 26% increase for **Deep Archive** storage.

There will be no changes to the existing premiums for IBM® Power® Virtual Server, third-party software, or network bandwidth.

PaaS global price changes

Effective 1 January 2024, all prices for IBM Cloud service will increase by 3% globally. This increase covers:

- IBM Cloud® Activity Tracker
- IBM Cloud® App ID
- IBM Cloud® Code Engine
- IBM Cloud® Continuous Delivery
- IBM Cloud® Data Engine
- IBM Cloud® Data Security Broker
- IBM Cloud® Databases for DataStax
- IBM Cloud® Databases for Elasticsearch
- IBM Cloud® Databases for EnterpriseDB

- IBM Cloud® Databases for etcd
- IBM Cloud® Databases for MongoDB
- IBM Cloud® Databases for MySQL
- IBM Cloud® Databases for PostgreSQL
- IBM Cloud® Databases for Redis
- IBM Cloud® Event Notifications
- IBM Cloud® Internet Services
- IBM Cloud® Kubernetes Service and Red Hat® OpenShift® on IBM Cloud®
- IBM Cloud® Log Analysis
- IBM Cloud® Messages for RabbitMQ
- IBM Cloud® Monitoring
- IBM Cloud Satellite®
- IBM Cloud® Secrets Manager
- IBM Cloud® Security and Compliance Center
- IBM Cloud® Security and Compliance Center Workload Protection
- IBM® Cloudant® for IBM Cloud®
- IBM® Event Streams for IBM Cloud®
- IBM® Key Protect for IBM Cloud®

Troubleshooting the IBM Cloud docs

Why isn't an IBM Cloud docs page available?

If you go to an IBM Cloud® docs page and the content doesn't display, try searching for the content instead to get you to the correct place.

What’s happening

You click a docs link, select a docs search result, or follow an existing bookmark to the IBM Cloud docs, and the following error is displayed.

> Sorry, this content isn't available.
>
> The information you're looking for might have moved. Try searching the IBM Cloud Docs to find related content.

Why it’s happening

The page might have been moved or removed entirely, or the link itself might be incorrect.

How to fix it

Search the IBM Cloud docs to find related information.

1. On the error page, click **Search in Docs** to go directly to the submitting feedback. In your feedback, include information about where you found the link to the missing page.

Why can't I access some IBM Cloud docs?

Some documentation in IBM Cloud® has limited visibility, where access is controlled by an allowlist. Only accounts that are given access to these docs can view them.

What’s happening

You try to view documentation in IBM Cloud, but the content doesn't load and the following error is displayed.

> Looks like you don't have access to this content.
>
> Click Console to open the IBM Cloud console and verify that you're in the correct account.

Why it’s happening

Access to documentation is granted by account rather than by user. The account that you're in doesn't have access to this content.

How to fix it

If you were previously able to access the content and you're a user in multiple IBM Cloud accounts, check that you're in the correct account.

a. In the IBM Cloud docs menu, click **Console** to open the IBM Cloud console.
b. Verify that the correct account is selected in the IBM Cloud console menu bar. If the account that's displayed in the console menu bar is not the account that has access, switch to the correct account by selecting the account from the drop-down menu.
c. Return to the IBM Cloud docs page that you were trying to access, and reload the page.

If you're still unable to access the docs, you might not be in an account that's on the allowlist. Contact your IBM Cloud representative or go to the Support Center for help.

Why don't I get results in the FAQ or tutorial library?

If you filter for content in IBM Cloud® docs FAQs or tutorials libraries and don't get results, try searching the IBM Cloud docs for related content.

What’s happening

You filter for content in the IBM Cloud docs FAQ library or tutorials library and see the following related message.

In the FAQ library:

> No FAQs found

> It seems we can't find what you're looking for. Try updating your filter selections.

In the tutorials library:

> No tutorials found
>
> Sorry, there are no tutorials that match your criteria. Try updating your filter selections.

Why it’s happening

There might not be any content with your selected filters, or a topic might not be set with the correct filter value.

How to fix it

Clear the selected filters, and try filtering the library again. If you still don't see any content, try searching all of the IBM Cloud docs for related content.

a. Go to the submitting feedback.

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America
2025-01-22