Support

1. IONOS CLOUD Support

24/7 enterprise-grade support for IONOS CLOUD customers. Start here for contact channels, agreements, and service docs.

Contact Support | Browse the Service Catalog | View price lists | FAQ

Quick Links

Contact Support

Phone and email channels. What to prepare before you call.

Service Catalog

Specs, availability, and limits across IONOS CLOUD services.

Support Agreements

Terms and commitments for IONOS CLOUD Support.

Prices

Current price lists by contracting entity and country.

FAQ

Get answers to the most commonly encountered questions across products.

Before you contact support

- Your **Support PIN** and the **email address** linked to your account.
- Affected resources (IDs, region, IPs, timestamps, error messages).
- Impact summary (what's broken, since when, and what changed).

> If you're using Data Center Designer (DCD), your Support PIN is part of your account security settings.

Browse by Topic

Compute Engine

Dedicated Core, vCPU servers, sizing, and platform limits.

Block Storage

HDD/SSD/DAS, performance classes, snapshots, and images.

Backup Services

Backup capabilities, supported platforms, and protection packs.

Object Storage (S3)

S3-compatible endpoints, features, and location availability.

Managed Kubernetes

Cluster capabilities, updates, and integration with IONOS CLOUD services.

Network File Storage

Managed NFS, scaling limits, bandwidth, and IOPS profiles.

Private Cloud Powered by VMware

VMware stack, scaling, add-ons, and available locations.

Public Cloud Locations

Regions, certifications, and current location availability.

Support Agreement

Support scope and commitments.

Ionos Cloud

Support agreement terms for IONOS CLOUD customers.

Pricing and contact

- Contact Information

- Prices

2. IONOS CLOUD SUPPORT AGREEMENT

2.1. IONOS CLOUD

IONOS CLOUD

This document describes the support services and partner obligations for IONOS CLOUD customers contracting under IONOS CLOUD SE.

1. Description of 24/7 IONOS CLOUD support service

IONOS CLOUD customers have access to a comprehensive range of support services designed to make the use of IONOS CLOUD services as efficient and trouble-free as possible. Here is a detailed list of the services available to these customers:

1.1 24/7 support access 365 days a year

Customers can always receive support from qualified personnel who are trained to provide fast and effective solutions.

The 24/7 IONOS CLOUD Support can be reached as follows:

For telephone numbers and email addresses by country, see **Contact Information**.

1.2 Transmission of information

Support tickets must meet the following criteria on both sides:

- The subject of the email should contain a unique ticket ID from the cloud customer's ticket system, if available,
- an automatic acknowledgement of receipt must be sent, and
- the processor of the record must be named
- Contact options (email/telephone)
- Information on the instances and service concerned
- Time stamp

Upon receipt of a support request, a trained system administrator will contact the cloud customer Support. Support requests are processed in compliance with the response times defined below.

The following response times apply to incoming support requests:

- Malfunction (service is not available or its use is restricted) < 1 hour
- Service or information request < 72 hours

Within the response times specified above, the cloud customer will receive a qualified statement. This includes either a statement on the completion of the process or, if the process has not yet been completed, an initial assessment of the support request and information on how to proceed. In the case of malfunctions, the cloud customer will additionally be informed about the extent of the malfunction and the expected duration of the malfunction elimination.

1.3 Responsibilities

The 24/7 IONOS CLOUD Support has no authorisation to log on to virtual resources of the Cloud customer, nor to make any changes.

1.4 Overview of support services for cloud customers at IONOS:

- Technical Support
    - Help with troubleshooting and solving technical problems in connection with the IONOS CLOUD services.
- Maintenance announcements and system updates
    - Regular information about planned maintenance work and system improvements
    - Announcements about required security updates

- Service requests
  - Processing specific requirements, such as Cloud Connect
  - Requests for additional services, including penetration tests, MS SQL licenses, and data upload services
- Help with account and invoice management
  - First point of contact for administrative matters relating to account management and billing
- Management of status page for incidents
  - Providing and updating information on the status page during incidents to inform customers about the current status of the services
- Provision of the **Reason for Outage** documents

2. Support agreement for Partner of IONOS

In addition to the general cloud customer services, IONOS partners have access to extended support levels and commitments tailored to their specific needs. The following additional services are available to them:

2.1 1st level support

The Partner undertakes to provide the following support and services to its end customers. The same applies for Distributors via their Resellers:

- Information and service requests regarding
  - Contract (contract details, general terms and conditions, GTCs etc.)
  - Customer master data
  - Accounting
- Subcontract Management
  - Management of individual contracts
  - Determination or modification of the resources
  - User administration
  - Deletion of unneeded resources and contracts
- Support of virtual machines, operating systems, and services outside the scope of IONOS Cloud services

    - Root cause analysis and remediation
- Monitor or subscribe of IONOS Cloud **status page**
- Execution or forwarding of maintenance announcements
- Forwarding technical problems and incidents regarding IONOS CLOUD products to 24/7 IONOS CLOUD Support (clause 1.2 lists the required information)
- Ticket responsibility and communication to the customer always remains in 1st level support

2.2 2nd level support

Optionally, a Partner in the Partner level "Platinum" can achieve 2nd level support status to ensure the best possible service delivery. For Distributors, 2nd level support status is mandatory. Partners are trained and certified accordingly in this regard.

In order to achieve and maintain 2nd level support status, the Partner undertakes to maintain the relevant certifications and to diagnose and qualify a fault to the best of its ability within the scope of the support. This includes the following aspects:

- Details about components involved
    - IPs, UUIDs, operating system, applications, and versions
- Performance measurement (tools specified by IONOS CLOUD)
    - Network (internal/external connection - bandwidth, PPS, latency)
    - Storage (bandwidth, IOPS, latency)
    - Compute (Stealtime, load, cpu/iowait)
- Detailed error pattern analysis
    - exact error description/derivation
    - Log messages
    - Extract Monitoring
    - Clear time information (date, time, time zone)
- Transmission of setup relevant details
    - Update status
    - Configuration details of the components involved in the malfunction
- Incidents and performance limitations must always be reported to IONOS CLOUD

2.3 3rd level support

The 3rd level support is provided exclusively by 24/7 IONOS CLOUD Support. The support services include the following:

- Troubleshooting with regard to the services offered by IONOS CLOUD
- Announcement and communication of maintenance announcements
- Service Requests, for example, Cloud Connect, Private Cross Connect, Reverse DNS, penetration testing, MS SQL licenses, data upload service

2.4 Ticket handling at IONOS

2.4.1 The Partner's end customers or the Distributor's Reseller will not be provided with a direct communication channel to the 24/7 IONOS CLOUD Support of IONOS CLOUD. The Partner shall provide 1st level support to its customers for the purpose of information and service requests, as well as support in case of malfunctions on the operating system, application, and configuration level of the virtually used resources of IONOS.

2.4.2 If restrictions are due to the services provided by IONOS CLOUD, these must be reported by a support employee of the Partner to the 24/7 IONOS CLOUD Support.

2.4.3 Escalations, for example, exceeded response times of 24/7 IONOS CLOUD Support response time, shall be communicated to the Head of 24/7 IONOS CLOUD Support either by telephone or by email directly by the Partner. The response time, as stated in clause 1.2, measured from the time the Partner Support employee has submitted a fault to IONOS CLOUD Support as described in clause 1.2. This escalation measure shall only be permitted after the regular support process described in clause 1.2 has been fully completed.

You can contact the IONOS CLOUD Support Management team using the contact details that are available in the partner portal.

IONOS may change the contact details of the Head of 24/7 IONOS CLOUD Support by providing five (5) business days' notice in text form.

2.4.4 If the information transmitted to 24/7 IONOS CLOUD Support in the context of a ticket changes (e.g. contact person, email addresses, telephone number, status of the 1st level support ticket), this change must be communicated to 24/7 IONOS CLOUD Support immediately. This is essential for a reliable exchange of information and thus a timely processing of support requests.

2.5 Authentication

To protect the customer and their virtual infrastructure, tickets can only be processed with prior authentication of the reporting party. For this purpose, we recommend creating an individual access plus support PIN in the Data Center Designer for each partner support employee, which the partner must do.

2.6 Responsibilities

The 24/7 IONOS CLOUD Support has no authorization to log on to virtual resources of the Partner or its customers, nor to make any changes.

2.7 Trainings

In order to ensure the best possible communication and cooperation with all partners, 24/7 IONOS CLOUD Support offers regular training sessions to Partner Support staff, either free of charge or at a charge. These include the following contents:

- Presentation of 24/7 IONOS CLOUD Support
- Incident management and request fulfilment processes of IONOS Cloud
- Interface 24/7 IONOS CLOUD Support

Furthermore, IONOS CLOUD offers separate training and certification, free of charge or at a charge, for the achievement of 2nd level support status for the Partners. The annual mandatory certifications ensure that a Platinum Partner or Distributor has at least three (3) trained administrators available at all times. This certification includes the aspects mentioned under clause 2.2 for the fastest and best possible root cause identification and elimination in active cooperation.

3. GENERAL INFORMATION

3.1. Glossary of Terms

A B C D E F G H I K L M N O P Q R S T U V W

A

Ansible

An automation tool that allows users to configure, deploy, and orchestrate advanced tasks, such as continuous deployments, or zero downtime rolling updates across IT infrastructure.

Application Load Balancer (ALB)

A pre-configured Virtual Data Center (VDC) element that operates at the application layer of the Open Systems Interconnection (OSI) model. It directs incoming application traffic to multiple targets to improve the scalability and availability of web applications.

ALB

See **#application-load-balancer-alb**Application Load Balancer.

Application Programming Interface (API)

A set of rules and protocols that enables different software applications to communicate and exchange data with each other in a standardized and structured manner.

API

See **#application-programming-interface-api**Application Programming Interface.

Authoritative Zone Transfer (AXFR)

A DNS protocol that copies an entire zone from a primary to a secondary server. It ensures consistent, synchronized DNS data and uses TCP for reliable transfer. IP restrictions or TSIG keys typically secure AXFR and are ideal for complete zone replication in authoritative DNS setups.

Availability Zone

You can assign a Zone to a Virtual Machine to guarantee distinct locations within a region from which the cloud services are provided. It provides distributed resources across multiple zones to build highly available fault-tolerant architectures.

B

Backup Unit Manager

An application that schedules, manages, and operates data backup processes on a computer, server, or network device. It is an integrated application that works on a client-server architecture for extracting data backup copies from a source computer or IT environment to a remote storage facility.

Beta

A stage in the product's lifecycle where it is made available to a limited number of users or a specific target audience for testing and gathering feedback. The beta phase typically follows the alpha phase and precedes the software's General Availability (GA) or official release.

Block Storage

A storage architecture in which data is stored as fixed-size raw blocks. Each block acts as an independent storage unit managed directly by the operating system or application, without a filesystem layer in between. Block storage is commonly used for databases, virtual machines, and applications that require low-latency, high-performance storage.

Bucket

A user-defined storage area in a cloud storage system is divided into folders and can be accessed or managed through APIs or user interfaces provided by the cloud storage provider.

Bring Your Own License (BYOL)

A model in which customers run software on cloud infrastructure using licenses they already own, rather than acquiring new ones from the cloud provider. Eligibility and transfer conditions are governed by the terms of the original software vendor.

BYOL

See **#bring-your-own-license-byol**Bring Your Own License.

C

Chef: Knife plugin

A configuration management tool that allows simplified configuration and maintenance of both servers and cloud provider environments through the use of common templates called recipes.

Command Line Interface (CLI)

It is a text-based interface that allows users to interact with a computer system or software application through commands entered through a command line interface, rather than through a graphical user interface (GUI).

CLI

See **#command-line-interface-cli**Command Line Interface.

Cloud Block Storage

Cloud Hard Disk Drive (HDD) and Solid State Drive (SSD) Block storage allow customers to use a dual-redundant storage system. Each block of storage the customer creates is stored

on two storage servers, providing active-active redundancy. For additional data protection, every storage server is based either on a hardware or software RAID system.

Cloud Computing

Cloud computing is the delivery of information technology services over a network, usually the internet. In the cloud computing model, infrastructure, data, and software are hosted by the vendor, and delivered to the user as a service.

Cloud Cubes

Virtual private service instances with shared resources. Cubes are a new addition to the Cloud product line, distinct from our established enterprise VM products. At the same time, they are fully integrated with the entire Cloud solution stack and can be used with advanced enterprise resources and features. They are also hosted on the same LANs as the Enterprise VM instances.

Cloud-init

A defacto multi-factor package that performs early initialization of a cloud instance. It is a set of community-maintained scripts that configures virtual machines based on information from the metadata service.

Cloud Storage

A model of computer storage in which data is stored in facilities (often multiple facilities) managed by a hosting company (cloud service provider) and is accessed remotely by the user over a network.

Compute Engine

A virtual machine (VM) or a collection of VMs hosted on the space.vars.ionos_cloud infrastructure. Compute Engine is further categorized into Dedicated Core server and vCPU Server, each containing a set of resources to support different workloads.

Configuration Management Tools

An automated approach to maintaining computer systems and software in a known, consistent, state.

Container

An abstract unit of software is an independent executable unit that contains everything needed to run an application. Containers have specific parameters and can run programs, work tasks, or specific tasks. The small size, high speed, and portability of containers are due to the absence of a guest operating system in each instance. This way, containers can use host OS components and resources.

Content Delivery Network (CDN)

A globally distributed network of servers designed to deliver web content quickly and efficiently by caching it close to the user's geographic location, reducing latency and improving load times. CDNs also enhance website security through **#ddos**DDoS protection and provide scalability to handle high traffic volumes.

CDN

See **#content-delivery-network-cdn**Content Delivery Network.

Core Network

A high availability core network at each location for the redundant connection of the product platform. All services provided by Cloud are connected to the internet through this core network.

Cross Connect

A physical, hardwired cable that provides a direct connection between two different termination locations within a data center. Cross connects enable colocation customers to establish high-performance, dedicated connectivity to one another within the data center – providing increased reliability and lower latency over network connections standardly used outside of a data center.

Container Storage Interface (CSI)

An attempt to standardize the interaction between Container orchestration engines and storage providers. This solution is being promoted by Kubernetes, Docker, and Mesos. It

means that the storage providers implement one addition, and Container orchestration platforms undertake to support the CSI interface.

CSI

See **#container-storage-interface-csi**Container Storage Interface.

Cube

A three-dimensional (3D) (or higher) range of values that are generally used to explain the time sequence of an image's data. It is a data abstraction to evaluate aggregated data from a variety of viewpoints.

D

Database as a Service (DBaaS)

A cloud computing service that allows users access to, and the use of, a cloud database system. Users can access this without purchasing and setting up their hardware, installing their database software, or managing the database themselves.

DBaaS

See **#database-as-a-service-dbaas**Database as a Service

Database Cluster

A collection of databases that is managed by a single instance of a running database server. After initialization, a database cluster will contain a database named Postgres, which is meant as a default database for use by utilities, users and third-party applications.

Data Center Designer (DCD)

A unique graphical tool for creating and managing Virtual Data Centers (VDC) in the cloud; configuration is intuitive and straightforward with a JavaScript-based graphical user interface.

DCD

See **#data-center-designer-dcd**Data Center Designer.

Distributed Denial of Service (DDoS)

It is a type of cyber attack in which multiple compromised computers or devices, known as a botnet, are used to flood a target system, or network with a massive amount of traffic, or requests. The goal of a DDoS attack is to overwhelm the target's resources, such as bandwidth, processing power, or memory, rendering the system, or network unavailable to legitimate users.

DDoS

See **#distributed-denial-of-service-ddos**Distributed Denial of Service.

Dedicated Core Server

An Infrastructure-as-a-Service (IaaS) platform that provides on-demand access to integrated solutions that ensure your application runs smoothly during peak loads.

Digital Asset

An entity that exists in binary format and comes with a right to use. If there's no right to use, then it is not considered an asset. Common examples include images, audio files, spreadsheets, multimedia, email, websites, digital documents, and so on.

Distributed Network

A computer network system that distributes the programming, software, and data across multiple computers. These computers may be nearby or geographically dispersed, but they work and communicate together as one system.

Docker Machine

A command-line tool used to provision and manage Docker hosts (virtual machines or physical machines) running the Docker Engine. It simplifies the process of creating, configuring, and managing Dockerized environments on various platforms. For more information, see **https://docs.docker.com/**Docker Documentation.

Docker Machine Driver

The component that is responsible for creating and managing the underlying infrastructure where Docker hosts are provisioned. Docker Machine supports multiple drivers, each tailored to work with specific platforms or virtualization technologies.

Domain Name System (DNS)

A system used to convert a computer's hostname into an IP address on the internet. For example, if a computer needs to communicate with the web server example.net, your computer needs the IP address of the web server example.net.

DNS

See **#domain-name-system-dns**Domain Name System.

E

Early Access (EA)

A stage in software development or product release where a limited set of users or customers are granted early or pre-release access to a product or service. It allows these select individuals or organizations to use, test, and provide feedback on the product before its official start, or general availability.

EA

See **#early-access-ea**Early Access.

Elastic Computing

The ability to dynamically provision and de-provision computer processing, memory, and storage resources to meet changing demands without worrying about capacity planning, and engineering for peak usage.

European Union (EU) Legal Sovereignty

The principle that data is subject only to the laws and legal protections of EU member states. In a cloud context, this means the infrastructure provider, the data centers, and the governing contracts are all based in the EU, ensuring that foreign entities (such as those under the **#us-clarifying-lawful-overseas-use-of-data-cloud-act**US CLOUD Act) cannot legally bypass European privacy protections.

External Network

Depending on the location, different capacities for transmitting data to or from the internet are available for operating the Cloud service. Due to the direct connection between the data centers at the German locations, the upstream can be used across locations.

F

Failover

Failover is switching to a redundant or standby computer server, system, hardware component, or network during failure, or abnormal termination of the previously active application, server, system, hardware component, or network in a computer network. Failover and switchover are essentially the same operation, except that failover is automatic and usually operates without warning, while switchover requires human intervention.

Firewall

A network security device that monitors incoming and outgoing network traffic and permits or blocks data packets based on a set of security rules. Its purpose is to establish a barrier between your internal network and incoming traffic from external sources (such as the internet) to block malicious traffic like viruses and hackers.

Flow logs

A feature that allows you to capture data related to IPv4 network traffic flows. Flow logs can be enabled for any network interface of a Virtual Machine (VM) instance and Network Load Balancer, and the public interfaces of the Network Address Translation (NAT) Gateway.

Forwarding rules

The configuration settings dictate how network traffic is forwarded from a source to a destination in the context of network devices, such as routers or switches. These rules determine the routing path and actions taken on incoming packets.

File Transfer Protocol (FTP)

It is a standard network protocol used for transferring files between a client and a server over a TCP/IP-based network, such as the internet.

FTP

See **#file-transfer-protocol-ftp**File Transfer Protocol.

G

General Availability (GA)

The stage in the lifecycle of a software product or service is when it is considered fully developed, stable, and ready for widespread use by the general public, or intended audience. It signifies that the product has completed its testing, bug fixing, and refinement phases, and is deemed suitable for production environments.

GA

See **#general-availability-ga**General Availability.

Gigabyte

We follow the binary interpretation of the gigabyte (also called gibibyte, 'GiB'), where a gigabyte equals 1024 megabytes, a megabyte equals 1024 kilobytes, and a kilobyte equals 1024 bytes. This defines one gigabyte as 1,073,741,824 bytes.

H

Hard Disk Drive (HDD)

It is a non-volatile storage device used for storing and retrieving digital data. HDDs use magnetic storage technology to store data. Data is read from and written to the platters

using read/write heads that move across the surface of the disks.

HDD

See **#hard-disk-drive-hdd**Hard Disk Drive.

Hot Plug

Hot swapping or hot insertion refers to the ability to add or remove hardware components from a system while it is powered on or in operation, without a system shut down or a restart. This capability allows for convenient and seamless replacement or addition of devices without disrupting the overall system.

Host Machine

A piece of physical hardware that hosts virtual machines.

Hybrid Cloud

A cloud computing environment that is comprised of a mix of private cloud, public cloud, and on-premises solutions. In a Hybrid Cloud, private and public cloud infrastructures remain distinct from one another but are bound together by technology that allows data and services portability between them.

I

IP Address

A unique address that identifies a device on the internet or a local network. IP stands for "Internet Protocol," which is the set of rules governing the data format sent from the internet or local network.

Incremental Zone Transfer (IXFR)

A DNS protocol that transfers only the changes (deltas) made to a zone since the last update rather than the entire zone file. It improves efficiency and reduces bandwidth compared to AXFR. IXFR requires both servers to support it and is commonly used to keep secondary DNS zones up to date with minimal data transfer.

Infrastructure as a service (IaaS)

It is a cloud computing model that provides virtualized computing resources over the internet. With IaaS, users can access and manage fundamental IT infrastructure components, such as virtual machines, storage, networks, and operating systems, as scalable, and on-demand services.

IaaS

See **#infrastructure-as-a-service-iaas**Infrastructure as a service.

Image

A copy of the entire state of a computer system, stored in some non-volatile form such as a file. A system is said to be capable of using system images if it can be shut down and later restored to the same state. In such cases, system images can be used for backup.

Indexes

A method of sorting data by creating keywords or a listing of the data.

Internal Network

An internal network, also known as a **#lan**LAN, refers to a network infrastructure that is confined within a specific physical location, such as an office building, campus, or data center. It is designed to facilitate communication and data sharing among devices and resources within the defined network boundary.

Internet Protocol (IP) Address

A unique address that identifies a device on the internet or a local network. They are an essential part of the internet's infrastructure and follow either the IPv4 or IPv6 addressing scheme.

Internet Protocol version 4 (IPV4) address

A unique numerical identifier that is assigned to devices on a network. It is part of the underlying Internet Protocol suite and is used to identify and locate devices on a computer

network. IPv4 addresses are 32 bits long, represented as a series of four decimal numbers separated by periods (e.g., 192.168.0.1). Each decimal number, also known as an octet, represents 8 bits of the address.

Internet Protocol version 6 (IPV6) address

A numerical identifier that is assigned to devices on a network. It is the successor to IPv4 and provides an expanded address space to accommodate the increasing number of connected devices in the modern internet era. IPv6 addresses are 128 bits long, represented as eight groups of hexadecimal digits separated by colons (e.g., 2001:0db8:85a3 0000:8a2e:0370:7334). This longer address format allows for a significantly larger number of unique addresses compared to IPv4.

IONOS CLOUD Object Storage

Create buckets and store objects with this S3-compliant service. It allows you to store and retrieve large amounts of data in the form of objects.

K

Kubernetes

An open-source container orchestration platform that automates the deployment, scaling, and management of containerized applications. It was originally developed by Google and is now maintained by the Cloud Native Computing Foundation (CNCF). Kubernetes provides a framework for managing the lifecycle of containerized applications across a cluster of nodes.

Kubernetes Clusters

A set of nodes that run containerized applications. Containerizing applications package an app with its dependencies and some necessary services. They are more lightweight and flexible than virtual machines. Kubernetes Clusters streamline the development, deployment, and management of applications across diverse environments.

L

Live Vertical Scaling

A technology that permits you to scale the number of CPU cores and amount of RAM while the server is running, without having to restart it. Vertical scaling involves adding or removing resources such as CPU, memory, storage, or network capacity to a single server, or virtual machine.

Local Area Network (LAN)

A computer network that spans a limited geographical area, such as a home, office building, or campus is typically privately owned, and operated, providing high-speed, and low-latency connections between connected devices. It connects computers, devices, and resources within the defined area, allowing them to communicate, and share data.

LAN

See **#local-area-network-lan**Local Area Network.

Load Balancer

The process of distributing a set of tasks over a set of resources (computing units), to make their overall processing more efficient. Load balancing can optimize the response time and avoid unevenly overloading some compute nodes while other compute nodes are left idle.

M

Managed Kubernetes

A cloud-based service that provides a managed and simplified environment for deploying, managing, and scaling Kubernetes clusters. Kubernetes is an open-source container orchestration platform that automates the deployment, scaling, and management of containerized applications.

Managed NAT (Network Address Translation) Gateway

A service provided by cloud providers to enable outbound internet connectivity for resources within a private subnet of a virtual network. NAT allows multiple devices within a private network to share a single public IP address when accessing the internet.

Marketplace

An online catalogue hosted by a cloud provider where third-party software vendors can list and distribute their products and services. Customers can browse, evaluate, and deploy partner offerings directly through the platform, simplifying the discovery, and procurement of software solutions.

Media Access Control (MAC) address

A unique identifier that is assigned to network interfaces at the hardware level. It is a 48-bit address that is used to identify devices on a local network.

MAC

See **#media-access-control-mac-address**Media Access Control (MAC) address.

Monitoring as a Service (MaaS)

An infrastructure provisioning model where physical hardware, such as servers, storage devices, and networking equipment, is treated as a service. MaaS allows users to manage and deploy bare-metal servers or physical infrastructure resources through a cloud-like interface or API.

MaaS

See **#monitoring-as-a-service-maas**Monitoring as a Service.

Multi-Cloud

The concurrent use of separate cloud service providers for different infrastructure, platform, or software needs. A multi-cloud approach can help prevent vendor lock-in and may help an enterprise deal with diverse workloads and partners.

Multi-Tenancy

A mode of operation for software in which multiple instances of one or many applications run in a shared environment. In a cloud computing model, pooled physical and virtual resources are dynamically assigned and reassigned to tenants according to consumer demand.

N

Network Address Translation (NAT)

A networking process that involves the modification of IP addresses in packet headers while they are being transmitted across a network. It is commonly used to enable communication between networks that use different IP address ranges.

NAT

See **#network-address-translation-nat**Network Address Translation.

Network Address Translation (NAT) Gateway

A networking component that enables communication between a private network and the internet. It acts as an intermediary device that translates IP addresses between the private network and the public internet. You can use a NAT gateway such that instances in a private subnet can connect to services outside your VPC but external services cannot initiate a connection with those instances.

Network File System (NFS)

A distributed file system protocol that allows remote file access and sharing over a network. NFS enables a client-server model, where a central server provides access to files and directories to remote clients.

NFS

See **#network-file-system-nfs**Network File System.

Network Load Balancer (NLB)

A pre-configured **#vdc**VDC element that provides connection-based layer 4 load balancing features and functionality. It improves the availability, scalability, and performance of network applications by efficiently distributing incoming traffic to multiple servers.

NLB

See **#network-load-balancer-nlb**Network Load Balancer.

Network Interface Card (NIC)

A component that provides networking capabilities for a computer. It may enable a wired connection (such as Ethernet) or a wireless connection (such as Wi-Fi) to a **#lan**LAN.

NIC

See **#network-interface-card-nic**Network Interface Card.

Network Security Groups

Network Security Groups (NSGs) are virtual firewalls that control the inbound and outbound traffic of Network Interface Cards (NICs) and Virtual Machines (VMs) in a Virtual Data Center (VDC). NSGs allow you to filter network traffic based on security rules you define for the associated member resources.

Node

A computer that connects to a blockchain network. 'Full' nodes store the entire distributed ledger from the first (genesis) block to the present and participate in blockchain consensus. 'Light' nodes are used by resource-constrained devices for application-specific purposes, such as reading specific blockchain data and submitting new transactions to full nodes for inclusion in the ledger.

Non-Volatile Memory Express (NVMe)

A standard hardware interface for Solid State Drives (SSDs) that uses the PCI Express (PCIe) bus. NVMe replaces traditional storage protocols, such as Serial ATA (SATA) and Small Computer System Interface (SCSI), which were originally designed for Hard Disk Drives (HDDs) and may not fully exploit the performance capabilities of SSDs.

O

Object Storage

A type of data storage architecture that stores data as objects rather than traditional file hierarchies or blocks. In Object Storage, each object is assigned a unique identifier and is stored in a flat address space, making it available, and scalable. Objects are stored in a flat address space and can be accessed through APIs using their unique identifiers or keys.

Object Storage Key Manager

A component or service that manages the keys used for accessing and securing objects stored in an Object Storage system. It provides a centralized platform for generating, storing, and managing encryption keys that are used to encrypt, and decrypt data stored in the Object Storage.

Optimization

The strategy of reducing the database system response time. Databases provide us with information stored in a hierarchical, related structure, which simplifies content extraction, and arrangement. There are plenty of options to choose from for businesses and companies.

P

Pay-As-You-Go (PAYG)

A flexible pricing model where you are billed only for the individual resources you use, for the time you use them, without any long-term contract.

Point-in-Time Recovery (PITR)

A data protection mechanism that allows you to restore a database or system to a specific point in time. It is commonly used in the context of databases but can also be applied to other types of systems. PITR works by capturing and storing incremental backups or transaction logs at regular intervals. These backups or logs contain the changes made to the data since the last full backup.

PITR

See **#point-in-time-recovery-pitr**Point-in-Time Recovery.

Platform as a Service (PaaS)

A cloud computing model that provides a managed environment for developing, running, and managing applications without the complexity of building and maintaining the underlying infrastructure. The cloud provider handles the operating system, runtime, middleware, and scaling, allowing teams to focus on application development.

PaaS

See **#platform-as-a-service-paas**Platform as a Service.

Provisioning

The process of setting up IT infrastructure. It can also refer to the steps required to manage access to data and resources and make them available to users and systems.

Provisioning is not the same thing as configuration, but they are both steps in the deployment process. Once something has been provisioned, the next step is configuration.

Public IPv4 Addresses

Every virtual network interface card that is connected to the internet is automatically assigned a public IPv4 address by DHCP. This IPv4 address is dynamic, meaning it can change while the Virtual Server is operational or during a restart.

Private IPv4 Addresses

In networks that are not connected to the internet, each virtual network interface card is automatically assigned a private IPv4 address. This is assigned by the DHCP service. These IPv4 addresses are assigned statically to the MAC addresses of the virtual network interface cards.

PTR Record

A DNS record that maps an IP address to its corresponding domain name, enabling **#reverse-dns-rdns**reverse DNS lookups. It is primarily used to verify an IP address's legitimacy, often in email validation scenarios.

Puppet

An open-source configuration management and automation tool that helps system administrators and DevOps teams manage and automate the deployment and configuration of software and infrastructure across a network of computers. It provides a declarative language and a framework for defining and enforcing the desired state of systems, also known as infrastructure as code.

Q

Query

A request for a database's data so we can retrieve or manipulate it. It is a command issued to retrieve or manipulate data from a Database Management System (DBMS).

R

Remote Console

A terminal or workstation in a remote location that is used to monitor and control a local computer.

Reverse DNS (rDNS)

The process of querying the DNS to resolve an IP address back to its associated domain name through **#ptr-record**PTR records. It is commonly used for email verification and network troubleshooting, helping identify the network traffic source.

Routing Table

A data table stored in a router or a network host that lists the routes to particular network destinations, and sometimes, metrics (distances) associated with those routes. The routing table contains information about the topology of the network immediately around it. A routing table is also known as a Routing Information Base (RIB).

S

Scalability

The ability of a process, system, or framework to handle a growing workload. In other words, a scalable system is adaptable to increasing demands. The ability to scale on demand is one of the biggest advantages of cloud computing.

Scale-out Backup Repository (SOBR)

A data backup feature from Veeam Backup and Replication that supports horizontal scaling of data storage in multiple tiers. It consists of several backup or Object Storage repositories.

SOBR

See **#scale-out-backup-repository-sobr**Scale-out Backup Repository.

Server Name Indication (SNI)

An extension of the SSL/TLS protocol that enables a client, such as a browser, to specify the domain to connect in the initial handshake of an HTTPS session.

SNI

See **#server-name-indication-sni**Server Name Indication.

Software as a Service (SaaS)

A cloud delivery model in which software is hosted and managed by a provider and made available to users over the internet, typically through a web browser or API. Customers access the application on a subscription basis without managing the underlying infrastructure, platform, or runtime environment.

SaaS

See **#software-as-a-service-saas**Software as a Service.

Software Development Kit (SDK)

A toolbox that contains special libraries that can be used to manage your cloud servers by using your preferred development or coding language. They hold tools, libraries, relevant documents, sample code, processes, and guides for use on that particular platform.

SDK

See **#software-development-kit-sdk**Software Development Kit.

Server

A piece of computer hardware or software (computer program) that provides functionality for other programs or devices, called "clients". This architecture is called the client-server model. Servers can provide various functionalities, often called "services", such as sharing data or resources among multiple clients or performing computations for a client.

Snapshot

A snapshot is the state of a system at a particular point in time. It can refer to an actual copy of the state of a system or a capability provided by certain systems.

Solid-State Drive (SSD)

A drive that uses integrated circuit assemblies to store data persistently, typically using flash memory.

SSD

See **#solid-state-drive-ssd**Solid-State Drive.

Sovereign-as-a-Service (SaaS)

A cloud delivery model in which a provider hosts and operates software or infrastructure within a defined legal and regulatory boundary, ensuring that data residency, access controls, and operational processes comply with the sovereignty requirements of a specific jurisdiction or customer agreement.

SaaS (Sovereign)

See **#sovereign-as-a-service-saas**Sovereign-as-a-Service.

Secure Shell (SSH) Keys

Used for easier authorization on various services. When creating storage based on Linux images, you can inject SSH keys into your VM to set up SSH access so that you can use SSH for secure communication.

Server Message Block (SMB)

A network file-sharing protocol that allows applications and users to read, create, and update files on remote servers.

Start of Authority (SOA)

A DNS record that defines key administrative information about a DNS zone: it includes the primary name server, the email address of the primary contact, and timers for zone refresh, retry, and expiry. The SOA record also contains the serial number, which is crucial for zone transfers, as it indicates when a zone has been updated.

Storage

A process through which digital data is saved within a data storage device utilizing computing technology. Storage is a mechanism that enables a computer to retain data, either temporarily or permanently.

T

Targets

The compute resources, such as VM instances, containers, microservices, or appliances, to which the traffic is distributed for processing. Network Load Balancer backend serves registered targets using an IP address and a TCP port.

Targets can be added or removed and capacities scaled without disrupting the overall flow of connection requests.

Terraform

An open-source Infrastructure as Code (IaC) tool to enable the provisioning, management, and deployment of infrastructure resources across various cloud providers, data centers, and other service providers in a declarative, and automated manner.

Token

A utility, resource, or asset value that can be bought, sold, or traded on an existing blockchain, like Ethereum. Tokens are often used in decentralized applications.

Two Factor Authentication (2FA)

A security mechanism that adds a layer of protection to user accounts or systems by requiring two forms of authentication to verify identity. The two factors involve a password or PIN and a mobile device or security token. By requiring both factors, 2FA mitigates the risk of unauthorized access, even if the user's password is compromised, enhancing the overall security of the system or application.

U

US Clarifying Lawful Overseas Use of Data (CLOUD) Act

The federal law that allows US law enforcement to compel US-based technology companies (through warrant or subpoena) to produce requested data, regardless of whether that data is stored in the US or abroad. European organizations subject to the GDPR must review the jurisdictional implications before proceeding.

V

Virtual CPU (vCPU)

A central CPU is associated with a Virtual Machine. It is also called a virtual processor.

vCPU

See **#virtual-cpu-vcpu**Virtual CPU

Virtual Data Center (VDC)

A collection of cloud resources used for creating an enterprise-grade IT infrastructure. VDC resources include the processors, memory, disk space, and networks from which virtual machines are built.

VDC

See **#virtual-data-center-vdc**Virtual Data Center.

VirtIO

VirtIO is a standardized interface for virtualized environments that provides drivers for various virtual devices, such as network interfaces, storage controllers, and other peripherals. Using VirtIO drivers with Kernel-based Virtual Machines (KVMs) offers several advantages, including improved performance and efficiency compared to emulated devices. VirtIO drivers are optimized for virtual environments and provide direct access to underlying hardware when available to enhance I/O performance and reduce overhead.

Virtual Machine (VM)

The virtualization or emulation of a computer system is based on computer architectures and provides the functionality of a physical computer. Virtual machines are based on computer architectures and provide the functionality of a physical computer.

VM

See **#virtual-machine-vm**Virtual Machine

Virtual Network

A software-defined network that allows you to create isolated virtual networks within a cloud or virtualized environment. It provides network connectivity and segmentation between different resources and allows them to communicate securely with each other.

Virtual Server

Scalable instances with dedicated resource functionality.

W

Write Ahead Logs (WAL)

A technique used in database systems to ensure data durability and consistency. It is a sequential log of changes that is written before the corresponding data modifications are applied to the database. WAL provides a reliable and efficient mechanism for data recovery during a system crash or failure.

WAL

See **#write-ahead-logs-wal**Write Ahead Logs.

Agent Instructions: Querying This Documentation

If you need additional information that is not directly available in this page, you can query the documentation dynamically by asking a question.

Perform an HTTP GET request on the current page URL with the `ask` query parameter:

```
GET https://docs.ionos.com/cloud/support/general-information/glossary-of-terms.md?ask=<qu
```

The question should be specific, self-contained, and written in natural language.
The response will contain a direct answer to the question and relevant excerpts and sources from the documentation.

Use this mechanism when the answer is not explicitly present in the current page, you need clarification or additional context, or you want to retrieve related documentation sections.

Service Catalog

Service Catalog is a central source of information about the services IONOS CLOUD offers to its customers.

Scope of Validity

This document offers details on the worldwide provision and operation of all services provided by IONOS CLOUD It is valid for IONOS CLOUD SE and its national affiliates.

The IONOS CLOUD platform

IONOS CLOUD provides a comprehensive, enterprise-grade cloud computing platform offering both Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) solutions. The portfolio encompasses scalable virtual and dedicated compute resources, high-performance block and object storage, and advanced software-defined networking, alongside fully managed services for Kubernetes, databases, and Artificial Intelligence workloads. Resources are provisioned on an on-demand basis, with billing calculated according to the specific consumption of compute, storage, and network capacity as defined in the current price list.

Users maintain full control over the provisioning and decommissioning of resources. IONOS CLOUD provides interfaces for this purpose, allowing customers to control the resources in a flexible manner. The following interfaces are currently available:

- Data Center Designer (DCD)
- Cloud API

IONOS CLOUD model

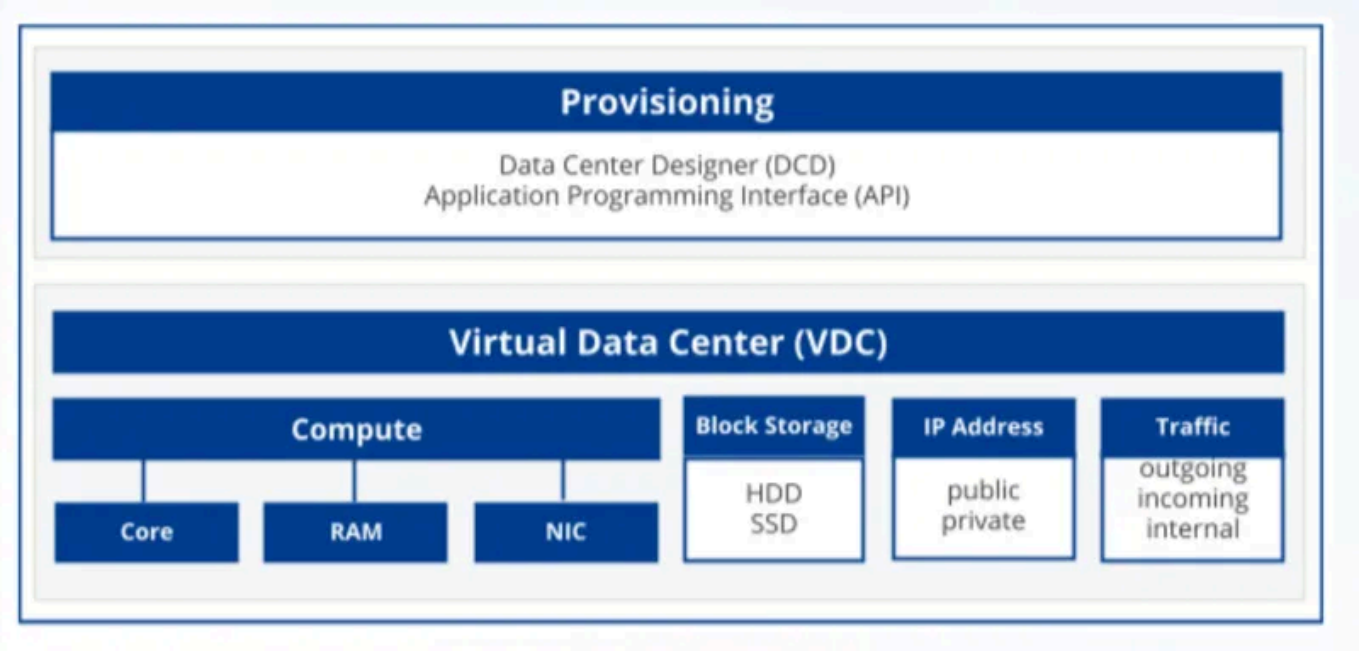

Data Center Designer

IONOS CLOUD provides the customer with access to a personalized web application called the “Data Center Designer” (DCD). The DCD can be accessed via modern Internet browsers. Specifically, the DCD allows the customer to both control and manage the services or sub-services provided by IONOS CLOUD, including:

- Creating, editing, and deleting virtual data centers
- Creating, (re-)starting, stopping, and deleting virtual servers, including optional storages
- Configuring/modifying existing virtual servers, including optional storages
- Creating, editing, and deleting snapshots
- Uploading, editing, using, and deleting private images
- Reserving and managing static public IP addresses
- Creating and managing private and public LANs, including firewall setups
- Creating and managing SSH keys
- Management of integrated Cloud services (For example, IONOS CLOUD Object Storage, Managed Kubernetes, Backup Service)

Virtual Data Center

On the IONOS CLOUD platform, the customer can create so-called “Virtual Data Centers” (VDC). A VDC is a repository for all infrastructure resources ordered by the customer. Access to the resources in a VDC – similar to operating a physical data center – is only possible via

a corresponding network or internet connection. Within a VDC, the IONOS CLOUD software allows for the distribution of various resources to different availability zones.

IONOS CLOUD provides the customer with the flexibility to change the ownership of the VDC. Any IONOS CLOUD customer who is a billing contract owner has the capability of transferring the ownership of a VDC created under the scope of his account with all related rights and responsibilities to any other customer having a billing contract with IONOS CLOUD. To change the ownership of his VDC the customer is requested to contact the 24/7 Enterprise Level Support.

Multi-User Management

Account Types

The authentication on the Data Center Designer requires that an account is assigned at least one username and one password. There are three distinct types of accounts:

- **Contract Owner** – This account is created automatically for the user who initially registered with IONOS CLOUD. Only one "Contract Owner" account can exist per contract made with IONOS CLOUD. A “Contract Owner” is authorized to fully access all resources, create and delete “User” accounts, and assign an “Administrator” role to them.
- **Administrator** - This role has the same privileges as the "Contract Owner", except this account type is restricted from changing the contract payment method. “Administrators” can assign “Administrator” roles to “User” accounts. It is possible to revoke the "Administrator" role after it has been assigned.
- **User** - This is the most basic account type. “Contract Owner” and "Administrator" account types can create or delete an unlimited number of accounts of the "User" type. This account type can be upgraded to the "Administrator" role and assigned specific privileges.

Resource Authorization

Multi-user management controls access to the following resource types:

- Virtual Data Centers
- Snapshots
- Images

- IP Blocks
- Backup Units
- Kubernetes Clusters

The assigning of rights is based on Groups. A Group contains one or more “User” accounts. A “User” account can be a member of several Groups. “Contract Owner” or “Administrator” accounts are exempt from group management, as they have access to all contracted resources.

Multi-user management distinguishes between the following authorizations for resources:

- **Read** – The resource and the objects they contain are visible. The attributes of objects can be displayed. The resource and the object it contains cannot be changed, nor can additional objects be added. The read authorization is implicit as soon as a group is assigned to a resource.
- **Edit** – The resource and the objects they contain are visible. The attributes of objects can be displayed and changed. The resource and the objects they contain can be deleted.
- **Sharing** – Authorizations for access to the resource can be changed.

Group Rights

The following group rights can be configured per group:

- **Create Data Center:** Create new virtual data centers.
- **Create Snapshots:** create snapshots of storages for which the group members hold at least “read” authorization.
- **Reserve IP Blocks:** reserve additional IPv4 blocks or give back IPv4 blocks, which are available to the group members via the “edit” authorization.
- **Create Internet Access:** Allows provisioning of public LANs inside a virtual data center.
- **Use Object Storage:** Access IONOS CLOUD Object Storage.
- **Create Backup Units:** Create a new Backup Unit account for backup agent registration and activation.
- **Create Kubernetes Clusters:** Create new Kubernetes Clusters.
- **Access Activity Log:** View Activity Logs for the entire contract.

Two-Factor Authentication

For every account that is configured for access to the IONOS CLOUD DCD (“Contract Owner”, "Administrator" or “User”), the use of 2-Factor Authentication can optionally be configured with a one-time password pursuant to RFC 6238 TOTP. 2-Factor Authentication provides increased security during the login process. In addition to their user name and password, when they log in, they are required to provide a code, which is generated using a special application (an “authenticator”).

Each account type can activate or deactivate this option in the DCD > Account management > Security for the respective account.

“Contract Owner” type accounts can set Two-Factor Authentication as a mandatory requirement for "Administrator" and “User” type accounts.

Remote Console

Integrated KVM Console: A fully interactive, browser-based HTML5 console providing direct out-of-band access to a server's monitor, keyboard, and mouse. This ensures secure access to the instance even if external network connectivity (SSH/RDP) is disrupted, without requiring additional plugins or client-side software.

IONOS CLOUD APIs

IONOS CLOUD provides a comprehensive suite of APIs that enable full programmatic control over the infrastructure. These APIs allow customers to automate the provisioning, configuration, and management of all resources—including Compute, Storage, Networking, and Managed Services—and integrate them directly into existing CI/CD pipelines and third-party tools.

Official Documentation & Specifications: For the complete, up-to-date list of available APIs, OpenAPI specifications, and technical endpoints, please refer to **IONOS CLOUD API Overview .**

AI

AI Model Hub

The AI Model Hub is a fully managed, serverless inference platform that provides secure API access to a curated selection of open-source Foundation Models. It enables the integration of generative AI capabilities—including text generation, code completion, and image creation—into applications without the need for customers to provision, manage, or scale the underlying physical GPU infrastructure.

Key Service Capabilities:

- **Foundation Model Access:** On-demand access to state-of-the-art Large Language Models (LLMs) and Text-to-Image models.
- **Retrieval Augmented Generation (RAG):** Integrated support for vector databases and text embeddings, allowing users to index proprietary document collections and ground model responses in their own private data.
- **API Compatibility:** Provides a fully OpenAI-compatible API endpoint, ensuring seamless interoperability with existing AI SDKs, libraries, and toolchains.
- **Usage-Based Scaling:** The platform operates on a serverless architecture with transparent rate limiting (Base and Burst limits) and billing based on consumed tokens or generated images.

Hosting Region: Germany (Berlin)

All data processing and inference occur exclusively within this region.

Service Availability: Global

The API endpoints are available from any location via the public internet.

Backup & Storage

Backup Service

The Backup Service is a fully integrated backup solution powered by Acronis technology. It enables customers to perform efficient backups of virtual machines, physical servers, and applications running within the IONOS CLOUD or on-premises.

Backup data is stored in ISO 27001 certified data centers, ensuring strict data sovereignty and security. Customers can choose to store data in the fully managed Backup Service repositories or target their own IONOS CLOUD Object Storage buckets. The infrastructure

leverages high-bandwidth internal connectivity (up to 10 GBit/s) to ensure short backup windows and seamless data throughput for both incremental and full system restores.

Features of the fully integrated backup function:

- Encrypted data storage in ISO-certified data centers in Germany
- Comprehensive image backup (full backup) or incremental backup
- Standard workload support
- Incremental and differential backups
- Basic scheduling
- Rapid disaster recovery (One-click recovery) and complete data restore
- Data migration
- Continuous data protection
- Backup support for Microsoft SQL
- Advanced database & cluster support: Application-aware and cluster-aware backups for Microsoft SQL Server (AAG), Microsoft Exchange (DAG), Oracle DB, SAP HANA, MariaDB, and MySQL.
- Data protection map and compliance reporting
- Off-host data processing
- Advanced virtualization tools
- Remote operations with bootable media

The Backup Service allows the customer to perform a quick and efficient backup of data from applications and any images the customer is using. The customer can also backup data from applications that run on-premises or in private Clouds in commonly used virtualized environments such as VMware and Hyper-V.

The fully integrated backup system supports the following platforms:

- Windows Server and Desktop operating systems
- Linux
- Mac
- Hyper-V and other Hypervisor solutions
- Workstations, physical and virtual servers

The following advanced protection packs can be enabled for each workload (a virtual machine, a server, or a workstation):

> **Note:** The Advanced Backup Feature Pack has been discontinued. All capabilities previously included in that pack — such as continuous data protection, SQL and cluster backup support, and off-host data processing — are now included in the base Backup Service at no additional cost.

- **Direct Backup to Public Cloud:** You can store backup data directly in public cloud storage instead of the default managed Backup Service repositories.
  - Backup directly to public cloud storage. For example, use our IONOS CLOUD Object Storage.
- **Advanced Security:** Protects your workloads continuously from all malware threats.
  - Antivirus and antimalware protection with local signature-based detection (with real-time protection)
  - Exploit prevention
  - URL filtering
  - Endpoint firewall management
  - Forensic backup, scan backups for malware, safe recovery, corporate allowlist
  - Smart protection plans (integration with CPOC alerts)
  - Centralized backup scanning for malware
  - Remote wipe
  - Microsoft Defender Antivirus
  - Microsoft Security Essentials
- **Advanced Management:** Allows you to patch vulnerabilities on the protected workloads.
  - Patch management
  - Disk health
  - Software inventory
  - Fail-safe patching
  - Cyber Scripting

- Remote assistance
- File transfer and sharing
- Selecting a session to connect
- Observing workloads in multi-view
- Connection modes: control, observe, and curtain
- Connection via the Quick Assist application
- Remote connection protocols: NEAR and Screen Sharing
- Session recording for NEAR connections
- Screenshot transmission
- Session history report – 24 monitors
- Threshold-based monitoring
- Anomaly-based monitoring

Block Storage

Block Storage allows the customer to use a dual-redundant storage system. Each block storage created by the customer is stored on two storage servers, providing active-active redundancy. For additional data protection, every storage server is based either on a hardware RAID system or on a software RAID system.

Direct Attached Storage (DAS) Block Storage based on Non-Volatile Memory Express (NVMe) are single-redundant storage systems. As this storage is installed directly into the physical server hosting the virtual machine, the storage volume is not stored across two servers. However, every DAS volume is covered by a software RAID system.

Access to the HDD and SSD volumes requested by the customer is achieved via the internal InfiniBand (RDMA) network. DAS volumes are connected to the mainboard of the server and benefit from fast peripheral component interconnect express (PCI express) bus performance.

For Solid State Drive volumes, IONOS CLOUD offers two performance classes that can be selected at the time of ordering the volume. SSD Premium is optimized for high performance while SSD Standard is recommended for fast data access with general-purpose scenarios.

HDD and DAS volumes deliver a static performance profile independent of the volume size. In comparison, SSD volumes deliver higher performance depending on the volume size and get

capped at a specific size.

Drive Type	Hard Disk Drive (HDD)
Use	Shared
Minimum and maximum size	1 GiB – 4 TiB per volume *(up to 24 HDD per VM supported)
Read / write speed, sequential	200 MB/s at 1 MiB block size
Read / write speed, random	1,100 IOPS at 4 KiB block size

**Larger volumes available on request.*

Drive Type	Solid State Drive (SSD) - Premium
Use	Shared
Minimum and maximum size	1 GiB – 4 TiB per volume * (up to 4 SSD per VM supported)
Read / write speed, sequential	1 MB/s per GiB at 1 MiB block size
Max. read / write speed, sequential	600 MB/s per VM at 1 MiB block size and min. 4 Cores, 4 GB RAM
Read speed, random	75 IOPS per GiB at 4 KiB block size
Max. read speed, random	45,000 IOPS per VM at 4 KiB block size and min. 4 Cores, 4 GB RAM
Write speed, random	50 IOPS per GiB at 4 KiB block size
Max. write speed, random	30,000 IOPS per VM at 4 KiB block size and min. 4 Cores, 4 GB RAM

**Larger volumes available on request.*

Drive Type	Solid State Drive (SSD) - Standard
Use	Shared
Minimum and maximum size	1 GiB – 4 TiB per volume * (up to 24 SSD per VM supported)
Read / write speed, sequential	0.5 MB/s per GiB at 1 MiB block size
Max. read / write speed, sequential	300 MB/s per VM at 1 MiB block size and min. 2 Cores, 2 GB RAM
Read speed, random	40 IOPS per GiB at 4 KiB block size
Max. read speed, random	24,000 IOPS per VM at 4 KiB block size and min. 2 Cores, 2 GB RAM
Write speed, random	30 IOPS per GiB at 4 KiB block size
Max. write speed, random	18,000 IOPS per VM at 4 KiB block size and min. 2 Cores, 2 GB RAM

**Larger volumes available on request.*

Drive Type	Direct Attached Storage (DAS) NVMe
Use	Shared
Minimum and maximum size	predefined per template 30 GiB - 640 GiB (1 DAS per Cubes VM)
Read / write speed, sequential	250 MB/s at 1 MiB block size
Read / write speed, random	5000 IOPS at 4 KiB block size
Bandwidth burst	500 MB/s at 1 MiB block sizefor 60 seconds
IOPS burst	10000 IOPS at 4 KiB block sizefor 60 seconds

Snapshot

IONOS CLOUD allows the customer to create so-called snapshots of individual block storages (HDD, SSD, DAS). A copy of each block storage can be accessed (and deleted) via DCD and Cloud API, and new block storage of any type can be created based on a snapshot. The provisioning speed is 50 MB/s.

Available Images

The following list provides an overview of the standardized images provided by IONOS CLOUD:

Open Source Linux

- Alma Linux
- Debian Linux
- Rocky Linux
- Ubuntu Linux

Enterprise Linux

- SUSE Linux Enterprise Server (SLES)

> **Note:**
>
> - A Bring Your Own Subscription (BYOS) model is applicable for SLES, and it is available in the following regions in Germany:
>   - Berlin `de/txl`
>   - Frankfurt `de/fra` and `de/fra2`

- Red Hat Enterprise Linux

Microsoft Windows Server

- Microsoft Windows Server 2016
- Microsoft Windows Server 2019
- Microsoft Windows Server 2022
- Microsoft Windows Server 2025

Microsoft SQL Server

- Microsoft SQL Server 2019 (Web, Standard, Enterprise)
- Microsoft SQL Server 2022 (Web, Standard, Enterprise)
- Microsoft SQL Server 2025 (Standard, Enterprise)

> **Note:** MS SQL Enterprise is only available on request. For more information, see **Microsoft SQL Enterprise Server Images** .

New versions of the standardized images may be added and old versions will be removed when the vendor no longer supports them.

> **Note:** We reserve the right to add non-LTS and testing/beta versions. Please follow the vendor's recommendations and refrain from using them for production use cases.

Image Upload

IONOS CLOUD allows the customer to upload their own images to the infrastructure via upload servers. This procedure is to be completed individually for each data center location. IONOS CLOUD optionally offers transmission with secure transport (TLS). The uploading of HDD and CD-ROM/DVD-ROM images is supported. Specifically, the uploading of images in the following formats is supported:

CD-ROM / DVD-ROM

- `*.iso` — 9660 image file

HDD Images

- `*.vmdk` VMware HDD images
- *.vhd, *.vhdx HyperV HDD images
- *.cow, *.qcow, *.qcow2 Qemu HDD images
- *.raw binary HDD image
- *.vpc VirtualPC HDD image
- *.vdi VirtualBox HDD image

A dedicated upload server is available for each data center location. Images can be transmitted to the upload server encrypted via FTPS (FTP-TLS) or unencrypted via FTP.

The following upload servers are available:

- Berlin: `ftp-txl.ionos.com`
- Frankfurt: `ftp-fra.ionos.com` (de/fra) and `ftp-fra-2.ionos.com` (de/fra/2)

- Karlsruhe: `ftp-fkb.ionos.com`
- London: `ftp-lhr.ionos.com`
- Worcester: `ftp-bhx.ionos.com`
- Paris: `ftp-par.ionos.com`
- Logroño: `ftp-vit.ionos.com`
- Las Vegas: `ftp-las.ionos.com`
- Newark: `ftp-ewr.ionos.com`
- Lenexa: `ftp-mci.ionos.com`

Once the image has been transmitted to the upload server, the image will be converted into the internal image format of IONOS CLOUD. The user will be informed by email when the conversion process starts.

Once the conversion is complete, the image will be available for use in the DCD or Cloud API under the name by which it was transmitted to the upload server.

Data Upload Service

IONOS CLOUD offers customers the ability to transfer large amounts of data via a physically mailed data storage medium. This service supports a variety of data carrier interfaces like USB or SATA. To ensure data security, the data on the delivered data storage medium must be encrypted and have a total size of at least 1 TB.

All uploads are performed as a one-to-one copy to a volume and provided in the data center chosen by the customer. The customer can attach this volume to a Virtual Server of their choice in the chosen virtual data center.

After the upload is complete, the data storage medium will be returned to the customer. The data upload service can be requested by the 24/7 Enterprise Level Support.

Storage Availability Zones

Customers can assign availability zones to HDD and SSD storage volumes to secure data, improve reliability, and create high availability scenarios.

Note that the Storage Availability Zone for SSD in data center location Karlsruhe is not provided. DAS storage volumes do not support availability zones as they are installed on the physical compute servers directly.

Virtual storage volumes with different storage availability zones assigned operate on different physical resources. Availability zones can be assigned using the DCD or Cloud API.

Cloud-init

IONOS CLOUD offers Cloud-init support for all Linux images. For Windows images, no Cloud-init functionality is provided. The feature is activated in all locations. All public IONOS CLOUD Linux images support Cloud-init. For private images, it is the customer's responsibility to make sure that their own images support Cloud-init.

Currently, IONOS CLOUD supports injecting user data.

Network File Storage

Network File Storage is a managed service that provides shared file storage to multiple virtual machines of the virtual datacenter using the NFS protocol.

The product is based on the ZFS open-source file system. It can detect and correct errors while in use without the need for a dedicated file system checker, making it suitable for mission-critical applications that require high availability.

Network File Storage uses 2 virtual machines in Active-Passive mode to provide high availability. Access to the SSD volumes is achieved via the internal InfiniBand (RDMA) network.

Storage is based on the SSD Standard performance class recommended for fast data access with general-purpose scenarios. Data is stored on two storage servers, providing active-active redundancy. For additional data protection, every storage server is based either on a hardware RAID system or on a software RAID system.

Network File Storage provides access using the NFS v4.2 protocol. It allows standard Linux clients to read and write directly to storage, scaling performance linearly for both IOPS and throughput, maximizing the limits of storage and network infrastructures:

- Reduced protocol overhead with compound operations and caching.
- Efficient file operations with minimal server interaction.
- Enhanced performance with multiple parallel network connections.
- Detailed file-level access and performance telemetry.

Drive Type	SSD - Standard
Usage mode	Shared
Minimum and maximum cluster size	2 TiB – 42 TiB
Maximum bandwidth	300 MB/s per 4 TiB with a 1 MiB block size
Maximum Read IOPS	24.000 IOPS per 4 TiB with a 4 KiB block size
Maximum Write IOPS	18.000 IOPS per 4 TiB with a 4 KiB block size

IONOS CLOUD Object Storage

IONOS CLOUD Object Storage is a secure, scalable storage solution that offers high data availability and performance. The product adheres to the S3 API standards, enabling the storage of vast amounts of unstructured data and seamless integration into S3-compatible applications and infrastructures.

It supports the following features:

- Versioning
- Logging (available for user-owned buckets)
- Server-side encryption with IONOS CLOUD-managed keys (SSE-S3)
- Server-side encryption with customer-provided keys (SSE-C)
- Access management via Bucket Policies and Access Control Lists (ACLs)
- Lifecycle management
- Static website hosting
- Cross-Region Replication (available for user-owned buckets)
- Object Lock with Governance and Compliance modes
- Bucket and object tagging
- Public Access Block
- Access from a private LAN to Object Storage endpoints using a Managed Network Load Balancer

IONOS CLOUD Object Storage is included with every contract, with no need for additional registration or activation. Through a user-friendly graphical interface and standard S3-compatible Object Storage clients, customers can efficiently manage their objects and

configure access controls using Bucket Policies in accordance with the Object Storage standard.

Currently, IONOS CLOUD Object Storage is available in the following locations, with more locations to be added soon.

Location	IONOS CLOUD Object Storage API Endpoint	Static Website	BSI IT Grundschutz	Bucket Type
Frankfurt (DE/FRA)	`s3.eu-central-1.ionoscloud.com`	`s3-website-de-central.profitbricks.com`	Compliant	User-owned
Berlin (DE/TXL)	`s3.eu-central-2.ionoscloud.com`	`s3-website-eu-central-2.ionoscloud.com`	Compliant	User-owned
Logroño (ES/VIT)	`s3.eu-south-2.ionoscloud.com`	`s3-website-eu-south-2.ionoscloud.com`	Compliant	User-owned
Frankfurt (DE/FRA/2)	`s3.eu-central-4.ionoscloud.com`	`s3-website.eu-central-4.ionoscloud.com`	Compliant	Contract-owned
Berlin (DE/TXL)	`s3.eu-central-3.ionoscloud.com`	`s3-website.eu-central-3.ionoscloud.com`	Compliant	Contract-owned
Lenexa (US/MCI)	`s3.us-central-1.ionoscloud.com`	`s3-website.us-central-1.ionoscloud.com`	Compliant	Contract-owned

Note: The hostname of the static website for location Frankfurt points to a ProfitBricks Object Storage installation which was the legal predecessor of IONOS CLOUD.

Compute Services

Compute Engine

Dedicated Core Server

These virtual machines run on dedicated CPU Cores. With Dedicated Core Servers you gain full access to the provisioned CPU resources,free from resource sharing with other virtual machines on the same physical host. This guarantees optimal performance, stability, reduced latency, and predictable performance. You can freely configure the number of cores and RAM required for your workloads, while choosing from the available CPU types available in your current VDC. Dedicated Core Servers can start from a storage volume, a CD-ROM, or a NIC.

Furthermore, it is also possible for a Dedicated Core Server to be configured with the help of advanced settings for the use of “Live Vertical Scaling” (LVS). LVS allows further resources to be added to your virtual machine while the operating system is in use. The scaling of resources without having to restart your virtual machine can be applied as follows:

- Upscaling: CPU, RAM, NICs, storage volumes,
- Downscaling: NICs, storage volumes.

For IONOS CLOUD provided public Images, LVS is activated by default. LVS capabilities on private images and snapshots can be changed before applying them to new instances. The Image Manager provides edit functionality to these properties. Linux supports all standard LVS functions, Windows server instances, however, only support Upscaling of CPU, NICs and storage volumes, and Downscaling of NICs at this time. LVS for RAM is possible starting from 1 GB RAM and in full increments of GB.

Possible configurations of a Virtual Server are presented in the table below:

AMD Processors

Components	Minimum	Maximum
Processor core (core)	1 Core	62 Cores
Random access memory (RAM)*	0.25 GB RAM	230 GB RAM
PCI Slots (NICs and storage volumes)	0 PCI Slots	24 PCI Slots
CD-ROM	0 CD-ROM	2 CD-ROMs

** Increment/decrement of 1 GB when LVS is activated, RAM expansion beyond the defined maximum size possible on request.*

Intel Processors

Components	Minimum	Maximum
Processor core (core)	1 Core	62 Cores
Random access memory (RAM)*	0.25 GB RAM	230 GB RAM
PCI Slots (NICs and storage volumes)	0 PCI Slots	24 PCI Slots
CD-ROM	0 CD-ROM	2 CD-ROMs

** Increment/decrement of 1 GB when LVS is activated, RAM expansion beyond the defined maximum size possible on request.*

Core

Generation	AMD EPYC Gen 3	AMD EPYC Gen 5	Intel Xeon Gen 5	Intel Xeon Gen 5	Intel Xeon Gen 5	Intel Xeon Gen 6
CPU Model	Milan	Turin	Haswell / Broadwell	Skylake	Ice Lake	Sierra Forest
Use	Exclusive	Exclusive	Exclusive	Exclusive	Exclusive	Exclusive
Clock frequency	2.0 GHz	2.6 GHz	2.0/2.1/2.4 GHz	2.0/2.1 GHz	2.0 GHz	2.4 GHz
Core Processing	Hyper-Threading	Hyper-Threading	Hyper-Threading	Hyper-Threading	Hyper-Threading	High-Throughput Cores

The different data center locations may be equipped with different CPU models. If the CPU model is listed below, but not displayed on the contract, please contact the IONOS CLOUD support team for further assistance.

CPU / Location	AMD EPYC Gen 3 (Milan)	AMD EPYC Gen 5 (Turin)	Intel Xeon Gen 5 (Haswell / Broadwell)	Intel Xeon Gen 5 (Skylake)	Intel Xeon Gen 5 (Ice Lake)	Intel Xeon Gen 6 (Sierra Forest)
Berlin (DE/TXL)	✓	–	–	✓	✓	✓
Frankfurt (DE/FRA)	✓	–	✓	✓	✓	–
Frankfurt (DE/FRA/2)	–	✓	–	–	–	✓
Karlsruhe (DE/FKB)	–	–	✓	–	–	–
London (GB/LHR)	–	–	–	✓	✓	–
Worcester (GB/BHX)	✓	–	–	–	✓	–
Paris (FR/PAR)	–	–	–	✓	✓	–
Logroño (ES/VIT)	–	–	–	✓	✓	–
Las Vegas (US/LAS)	✓	–	✓	–	–	–
Newark (US/EWR)	–	–	✓	–	✓	–
Lenexa (US/MCI)	✓	–	–	–	✓	–

Legacy CPU Models

The AMD Opteron CPU is classified as a legacy model and is not available for new provisionings. Its availability is limited to supporting existing workloads within the following locations:

- Karlsruhe (DE/FKB)
- Las Vegas (US/LAS)
- Newark (US/EWR)

AMD Core

IONOS CLOUD offers EPYC processors in selected locations, optimized for high performance within Cloud infrastructure.

The current processor design allows for a high number of cores within one host system. Therefore, virtual machines may have up to 62 AMD cores.

Intel Core

The Intel Xeon processors provided by IONOS CLOUD include both Gen 5 and Gen 6 models, each utilizing advanced multi-core technologies.

Intel Xeon 5 cores from the Haswell, Skylake, and Ice Lake families support hyper-threading, allowing each physical core to handle two simultaneous threads. This can improve performance by enabling more efficient parallel processing, provided the software supports multi-threading. Modern operating systems, including Windows and Linux, can fully leverage this feature.

The latest Gen 6 Intel Xeon processors, including Sierra Forest, use a multi-core architecture, where each physical core is presented as a single core within your virtual machine. This design ensures efficient and optimized resource allocation for demanding workloads in cloud environments.

Host Systems

Many host systems are kept ready at each location for operating virtual servers for the customers. Each host server is redundantly connected to the InfiniBand network. The host systems are assembled by the manufacturers based on our specifications and then delivered to the site.

vCPU Server

vCPU servers provide a good balance of compute resources, which are ideal for a wide range of applications. Unlike Dedicated Core Servers, which come with guaranteed dedicated resources, vCPU servers do not. This means that CPU resources are optimized by the hypervisor. The vCPU servers are well-suited for typical workloads that prioritize cost-efficiency and can accommodate variations in performance. Within the DCD or Cloud API, you can freely configure the ratio of vCPUs to RAM for your virtual machines. However, unlike the Dedicated Core Servers product, you do not have the option to choose the CPU type for your vCPU server.

Currently, vCPU Servers are available in:

- Berlin (DE/TXL)
- Frankfurt (DE/FRA)
- Frankfurt (DE/FRA/2)
- London (GB/LHR)
- Worcester (GB/BHX)
- Paris (FR/PAR)
- Logroño (ES/VIT)
- Newark (US/EWR)
- Lenexa (US/MCI)

vCPU Servers can start from a storage volume, a CD-ROM, or a NIC.

Additionally, vCPU Servers can use "Live Vertical Scaling" (LVS) for advanced configuration. LVS enables adding resources like vCPUs, RAM, NICs, and storage to a running Virtual Server without requiring a restart. This scaling process can be done to increase or decrease resources.

Limitations

Components	Minimum	Maximum
vCPU	1 vCPU	60 vCPUs
RAM	0,25 GB RAM	230 GB RAM*
NICs and storage	0 PCI connectors	24 PCI connectors
CD-ROM	0 CD-ROMs	2 CD-ROMs

- *Larger RAM sizes can be made available on request.*
- While provisioning the vCPU Server product, users cannot select the CPU Model through the DCD or Cloud API.
- Customer cannot auto-migrate from vCPU Server type to Dedicated Core Virtual Servers. A virtual machine recreation process is required.
- vCPU can be single or multi-threaded, depending on the underlying hardware.

Cubes

Cubes are a separate type of virtual machine. While Virtual Servers have the option to use Cores exclusively as 'Dedicated Cores', Cubes share them with other Cubes instances and expose virtual CPUs (vCPU). Still, these virtual machines are fully isolated and separated so that no data is available to any other virtual machine running on the same physical core.

In addition, Cubes are delivered with one NVMe storage that is directly attached to the physical server unit. This block storage device uses one of the PCI slots available by default.

Cubes are designed for cost optimization and workloads for which failover gets realized by the application and not the infrastructure. Cubes are currently rolled out to all European locations but may not be available to specific virtual data centers even if the feature is announced for availability in a particular location. Due to technical dependencies, Cubes may not be available for all legacy virtual data centers in the location Frankfurt. The product is available for newly created virtual data centers. Currently, Cubes are released in:

- Berlin (DE/TXL)
- Frankfurt (DE/FRA)
- Frankfurt (DE/FRA/2)
- London (GB/LHR)
- Worcester (GB/BHX)
- Paris (FR/PAR)
- Logroño (ES/VIT)
- Newark (US/EWR)
- Lenexa (US/MCI)

A Cubes instance consists of the following components:

- Virtual CPUs
- Memory (RAM)
- Network interface cards NIC (optional)
- Direct Attached NVMe Storage volume (mandatory)
- Block Storage volumes (optional)
- CD-ROMs (optional)

In comparison to Virtual Servers, Cubes get ordered by pre-defined instance size templates and cannot be configured in a fully flexible model.

List of available Cubes instances:

Name	vCPU	RAM	DAS Storage
Basic Cube XS	1	2 GB	60 GB
Basic Cube S	2	4 GB	120 GB
Basic Cube M	4	8 GB	240 GB
Basic Cube L	8	16 GB	480 GB
Basic Cube XL	16	32 GB	960 GB
Memory Cube S	2	8 GB	120 GB
Memory Cube M	4	16 GB	240 GB
Memory Cube L	8	32 GB	480 GB
Memory Cube XL	16	64 GB	960 GB

Cubes can start from any storage volume, a CD-ROM, or a NIC.

This type of instance does not support "Live Vertical Scaling" (LVS) of CPU or RAM even if it is enabled on an image. It is impossible to migrate to higher or lower tiers of Cubes.

LVS is limited to NICs and block storage volumes. Please note that the number of directly attached NVMe storage volumes is limited to 1 (one) and it cannot be expanded, deleted, removed from the Cubes instance, or migrated to any other instance. Attaching further block storage volumes must be of type HDD or SSD. The scaling of resources without having to restart a Virtual Server can be applied as follows:

- Upscaling: NICs, HDD/ SSD storage volumes
- Downscaling: NICs, HDD/ SSD storage volumes

Additional Services

IONOS CLOUD Cubes can be used inside a virtual data center with any other service provided in this location.

Host Systems

IONOS CLOUD operates different types of host systems based on AMD and Intel CPU architecture. All systems are configured to deliver the same performance. A specific CPU

type cannot be selected by the customer nor guaranteed by the IONOS CLOUD.

Cloud GPU VM

Cloud GPU VMs deliver high-end accelerated computing with dedicated NVIDIA GPUs, optimized for demanding AI and High-Performance Computing (HPC) workloads.

Cloud GPU VMs come with predefined templates. You cannot create new templates or modify the existing templates. Use Cloud API standard operations with template read access and volume management to deploy Cloud GPU VMs. Plan your storage requirements carefully and select the best template for your use case; otherwise, you must use additional volumes for your storage needs.

The product is available for Virtual Data Centers in the following locations:

- Frankfurt (DE/FRA/2)

The server sizing model allocates Cloud GPU VMs with corresponding dedicated CPU cores and RAM based on available host capacity. The architecture uses PCIe passthrough for direct hardware access and optimal performance.

You may choose between the following templates:

Name	GPU Model	GPUs	Dedicated vCPUs	RAM	Storage
H200 – S	H200 PCIe	1	15	267 GiB	1.000 GiB
H200 – M	H200 PCIe	2	30	534 GiB	1.536 GiB
H200 – L	H200 PCIe	4	60	1.068 GiB	2.048 GiB
H200 – XL	H200 PCIe	8	127	2.136 GiB	4.000 GiB

The templates can only be used with the Cloud GPU VMs. CPU and RAM allocate proportionally to the number of GPUs. Resources use dedicated cores with limited flexibility.

The first connected volume serves as the start volume, containing the operating system and required system files. Provision start volumes with adequate capacity at the initial Cloud GPU VM provisioning, because they use fixed sizing and cannot be detached or upscaled after deployment.

Known Limitations

- This section outlines important policies and infrastructure limitations for GPU Cloud VMs.

Instance lifecycle policy

- **Maintenance restarts:** While we prioritize planned maintenance windows, critical security patches, and urgent stability issues may require emergency maintenance without advance notification. Design your applications to handle unexpected restarts and configure services to restart automatically on start.
- **Maintenance Downtime:** GPU Cloud VMs do not support seamless background transfers during maintenance. Consequently, some planned infrastructure updates will require a temporary shutdown or restart of your instance.
- **Fixed resource scaling:** Size your servers appropriately at creation time. CPU core and RAM Upscaling or Downscaling are not currently supported after provisioning.

Management restrictions

- **API only support:** You can manage GPU Cloud VMs only via the Cloud API. You cannot use the Data Center Designer (DCD) or other interfaces for management.
- **Templates:** You cannot create new templates or modify existing templates.
- **Images:** The platform supports IONOS CLOUD Linux images only at start.
- **Start volume restrictions:** The first connected start volume has the following restrictions:
  - It cannot be detached from the VM.
  - It cannot be upscaled after deployment.
  - Size must be determined during the initial provisioning.
- **Resource scaling restrictions:** Size servers appropriately at creation time, because CPU core and RAM Upscaling or Downscaling are not currently supported. Plan your workload requirements carefully during the initial provisioning.
- **Volume Management:** Additional volumes support standard management operations, including scaling up, detaching, and attaching across the Cloud API interface.

Default Quotas

- By default, you can deploy a maximum of **1 (one) H200 – S** Cloud GPU VM instance per contract.

- To increase this limit, contact **IONOS CLOUD Support**.

Containers

Managed Kubernetes

The Managed Kubernetes is a managed service based on the open-source system for automating the deployment, scaling, and management of containerized applications. It groups Containers that make up an application into logical units for efficient management and discovery. It reduces the operational overhead of Container technology while maintaining development flexibility.

Managed Kubernetes facilitates the fully automated setup of Kubernetes Clusters. Several Clusters can be quickly and efficiently deployed, for example, to set up staging environments, and terminated when no longer needed. Kubernetes also significantly simplifies the automation of CI/CD pipelines in terms of testing and deployment.

Managed Kubernetes offers maximum transparency and control of the Kubernetes cluster. This includes:

- Fully automated setup of entire Kubernetes Clusters and Kubernetes node pools (with optional horizontal auto-scaling of nodes)
- Highly available and geo-redundant control plane
- Full cluster admin-level access to Kubernetes API
- Dedicated CPU and memory resource assignment
- Double redundant and persistent HDD/SSD storage
- Integration of Cloud services
- Regular security and version updates

Managed Kubernetes is free of charge. The customer pays only for the underlying IONOS CLOUD infrastructure that is actually needed.

Please note that IPv6 support has limitations for the IONOS CLOUD Managed Kubernetes service.

Private Container Registry

Private Container Registry is a universal repository manager that stores and manages custom Container images and other OCI-compliant artifacts. It can be used as part of CI/CD workflows for Container workloads in IONOS CLOUD Managed Kubernetes setups. Specifications for the Private Container Registry are as follows:

- Highly available service is managed, including any components on which it is built.
- Located in the region of Frankfurt am Main (DE) (further locations will follow).
- Support of the Docker Registry HTTP API V2.
- More than one repository per registry.
- Support of permanent and temporary access authentication tokens.
- Data encrypted when stored.
- Garbage Collection to release storage space.
- Vulnerability scanning of artifacts to identify security vulnerabilities.

Available locations: Frankfurt

Data Analytics

Event Streams for Apache Kafka

Fully managed Apache Kafka service delivering single-tenant clusters for real-time event streaming while IONOS CLOUD handles deployment, scaling, patching, and monitoring.

Key Features

- Provision Kafka clusters natively inside an IONOS Public Cloud VDC
- Fixed size profiles (XS – XL) with automatic fail-over
- TLS-encrypted endpoints; user & ACL management via API or Data Center Designer
- Optional private-LAN peering for internal workloads

Databases

The IONOS CLOUD Database as a Service (DBaaS) is a comprehensive web service designed to simplify the setup, operation, and scaling of databases within the IONOS CLOUD. This managed solution offers cost-efficient, resizable capacity for industry-standard databases, ensuring high availability, performance, and reliability. Hosted and integrated within the

IONOS CLOUD ecosystem, it alleviates common database administration tasks, providing an integration with other IONOS CLOUD services.

IONOS CLOUD DBaaS currently supports the following engines:

- In-Memory DB
- MariaDB
- MongoDB
- PostgreSQL

Every database engine in our DBaaS has a unique set of supported features designed to improve performance and functionality. Additionally, as each database engine evolves, new versions introduce distinct features, ensuring that database capabilities continuously improve and adapt to the latest technological advancements.

Network Services

Virtual Data Center (VDC) Networking

Virtual Network

IONOS CLOUD allows virtual entities to be equipped with network cards (“network interface cards”; NICs). Only by using these virtual network interface cards is it possible to connect multiple virtual entities together or to the Internet.

Parameter	Size	Performance
Throughput, internal	MTU 1,500	Up to 6 Gbps
Throughput, external	MTU 1,500	Up to 6 Gbps

The maximum external throughput may only be achieved with a corresponding upstream of the provider.

Compatibility

- The use of virtual MAC addresses or the changing of the MAC address of a network adapter is not supported. Among others, this limitation also applies to the use of CARP (Common Address Redundancy Protocol).
- Gratuitous ARP (RFC 826) is supported.

- Virtual Router Redundancy Protocol (VRRP) is supported based on gratuitous ARP. For VRRP to work, IP failover groups must be configured.

IP Address Management

IONOS CLOUD provides the customer with both IPv4 and IPv6 public IP addresses that, depending on the intended use, can be booked either permanently or for the duration for which a Virtual Server exists. Currently, only IPv4 addresses can be booked by the customer. These IP addresses provided by IONOS CLOUD are only needed if connections are to be established over the Internet. Internally, virtual machines can be freely networked. For this, IONOS CLOUD offers a DHCP server that allows or simplifies the assignment of IP addresses. However, one can establish one's own addressing scheme.

Public IPv4 Addresses

Every virtual network interface card that is connected to the Internet is automatically assigned a public IPv4 address by DHCP. This IPv4 address is dynamic, meaning it can change while the Virtual Server is operational or when restarted.

Customers can reserve static Public IPv4 Addresses for a fee. These reserved IPv4 addresses can be assigned to an Internet-connected virtual network interface card as primary or additional IP addresses.

Private IPv4 Addresses

In networks that are not connected to the Internet, each virtual network interface card is automatically assigned a private IPv4 address. This is assigned by the DHCP service. These IPv4 addresses are assigned statically to the MAC addresses of the virtual network interface cards.

The use of the IP address assignment can be enabled or disabled for each network interface card. Any Private IPv4 Addresses pursuant to RFC 1918 can be used in private networks.

Network address range	CIDR notation	Abbreviated CIDR notation	Number of addresses	Number of networks per network class (historical)
10.0.0.0 to 10.255.255.255	10.0.0.0/8	10/8	224 = 16.777.216	Class A: 1 private network with 16,777,216 addresses; 10.0.0.0/8
172.16.0.0 to 172.31.255.255	172.16.0.0/12	172.16/12	220 = 1.048.576	Class B: 16 private networks with 65,536 addresses; 172.16.0.0/16 to 172.31.0.0/16
192.168.0.0 to 192.168.255.255	192.168.0.0/16	192.168/16	216 = 65.536	Class C: 256 private networks with 256 addresses; 192.168.0.0/24 to 192.168.255.0/24

Public IPv6 Addresses

Every virtual data center is assigned a public /56 IPv6 CIDR block by default. Customers can choose to enable IPv6 in a LAN per their needs and a maximum of 256 IPv6 enabled LANs can be created per VDC. On enabling IPv6 in a LAN, the customer can either select a /64 IPv6 CIDR block from the /56 IPv6 CIDR block assigned to the VDC or have a /64 block automatically assigned to the LAN. Public IPv6 addresses are assigned to both private and public LANs.

Every connected virtual network interface card is then assigned a /80 IPv6 CIDR block and a single /128 IPv6 address either automatically, or the customer can also select both. They must though both be assigned from the /64 IPv6 CIDR block assigned to the corresponding LAN. The first public IPv6 address is assigned by DHCP and in total a maximum of 50 IPv6 addresses can be assigned per network interface. IPv6 addresses are static, meaning they remain assigned when a Virtual Server restarts.

DHCP

For every network interface of a Virtual Server, IONOS CLOUD provides an IP configuration via DHCP. In this context, the type of configuration distinguishes between whether the network

interface is connected with the public Internet or a private Ethernet.

Public Internet

The following parameters are provided for the configuration via DHCP:

- Public IPv4 address
- Network mask (255.255.255.255)
- Gateway address
- DNS server address
- MTU (1,500)

Similarly, DHCPv6 is supported for IPv6 public addresses

Private Networks

The following parameters are provided for the configuration via DHCP:

- Private IP address (for example, `10.0.0.1`)
- Network mask (255.255.255.0)
- MTU (1,500)

The DHCP server always uses the address A.B.C.1 in the class C network, which corresponds to the assigned IP address.

The configuration through DHCP can be optionally activated or deactivated via network interface (DCD, or Cloud API). The configuration via DHCP is activated for newly created network interfaces.

IP-failover

The IONOS CLOUD IP-failover feature helps to minimize packet loss for high availability or failover setups if one of the virtual machines experiences an outage. By setting up IP-failover groups for public traffic, customers can define the network interfaces of virtual servers that are part of a high availability setup.

“User” type accounts can create or edit IP-failover groups using only reserved IP addresses, for which they have been granted access. The IP-failover feature only provides provisioning of the same IP to multiple network interfaces from different Virtual Servers on the same LAN. It does not monitor the availability of the service to be accessed by the given IP. The

monitoring and GARP announcements to gateways must be made by the customer individually on each Virtual Server that is a member of an IP-failover setup.

IP failover groups cannot be created for IPv6 addresses and is an IPv4 only feature.

Connectivity & Hybrid Cloud

Cloud Connect

Cloud Connect enables the customer to create a direct and dedicated Layer-3 connection between their company network and their virtual data center (VDC). The customer can use Cloud Connect if both of the following conditions are fulfilled:

1. The connecting VDC is operated at the locations of Frankfurt (DE/FRA), Berlin (DE/TXL), or Las Vegas (US/LAS).
2. The customer has a dedicated line connection to the corresponding data center.

A connection can take place in different ways, for example, Dark-Fiber, MPLS, or Cross Connect. For this purpose, the customer can contract a telecommunications company to establish the connection.

The 24/7 Enterprise Level Support is available to assist with any questions about Cloud exchange and connection.

Managed NAT Gateway

In all locations, IONOS CLOUD provides a Managed Network Address Translation (NAT) Gateway. This service is exposing a Source NAT gateway which means it allows access from the virtual instance to the internet but blocks requests from the internet to the virtual infrastructure. This enables internet access to virtual machines without exposing them to the internet by a public interface. While being "hidden" from the internet and not being exposed to threats, the virtual machine still can initiate a connection to the customizable targets on the internet, e.g., to download new software updates or patches.

A Managed NAT Gateway requires the configuration of a reserved public IP address for the target configuration. The Managed NAT Gateway allows the configuration of multiple, individual NAT rules which can be applied to virtual machines being members of the listener LAN individually. These rules allow dedicated configuration of target subnets and port ranges which are explicitly allowed to be accessed by virtual machine instances.

The Managed NAT Gateway supports TCP, UDP, and ICMP protocols and up to six private networks per NAT Gateway. The number of Managed NAT Gateways per contract is limited to five (5) gateways. This limit can be adjusted by contacting the IONOS CLOUD Support Team.

The Managed NAT Gateway supports the recording of Flow Logs.

The Managed NAT Gateway operates in high-availability mode and service recovery is executed within seconds. The Managed NAT Gateway will be regularly maintained by IONOS CLOUD and updated with the latest software versions and new features. IONOS CLOUD reserves a weekly maintenance window that it can use for regular updates. It is scheduled every Monday between 02:00 - 04:00 AM local time of the data center in which the Managed NAT Gateway service is deployed. During maintenance, a service interruption of up to 5 seconds may occur.

Additional update deployments may be possible and carried out outside the maintenance window, for example, when urgent security patches are needed.

> **Note:** IPv6 support has limitations for the Managed NAT Gateway.

VPN Gateway

The fully managed VPN Gateway service provides secure and scalable connectivity, enabling encrypted communication between your IONOS CLOUD resources in a VDC and remote networks (on-premises, multi-cloud, private LANs in other VDCs, etc.).

A VPN Gateway requires configuring a reserved public IP address for the gateway configuration. VPN Gateway supports IPSec and WireGuard VPN protocols and allows the configuration of multiple IPSec tunnels or WireGuard peers for remote connectivity.

> **Note:** IPv6 support has limitations for the VPN Gateway.

Traffic Management & Performance

Classic Load Balancing

IONOS CLOUD offers the customer the function of a Classic Load Balancer for public traffic within their infrastructure. This load balancer distributes the incoming network traffic according to an ECMP algorithm on the servers configured behind the load balancer. The

Classic Load Balancer is for basic balancing scenarios and does not provide granular configuration or health checks.

Setting	Value
Throughput	50 Mbps
Max. open connections	40,000 connections
Max. backend servers	100

Managed Network Load Balancer

IONOS CLOUD offers a Managed Network Load Balancer (NLB) that is balancing layer 4/ TCP-based network traffic. This service is available in all locations.

Network Load Balancers can be provisioned as a private and a public load balancer. A public load balancer requires the configuration of a reserved public IP address for the target configuration. The network load balancer allows the configuration of multiple, individual load balancer rules which can be applied to virtual machines being members of the listener LAN.

The Network Load Balancers support multiple load balancing algorithms.

- Round Robin
- Least Connection
- Random
- Source IP

Furthermore, it offers options to specify health parameters to include or exclude nodes from the balancing configuration and manually remove listener targets from the load balancer (for example, when the node is in maintenance).

The number of Managed Network Load Balancers per contract is limited to five (5) NLBs. This limit can be adjusted by contacting the IONOS CLOUD Support Team.

The Managed Network Load Balancer supports the recording of Flow Logs.

The Managed Network Load Balancer operates in high-availability mode and service recovery is executed within seconds. The Managed Network Load Balancer will be regularly maintained by IONOS CLOUD and updated with the latest software versions and new features. IONOS CLOUD reserves a weekly maintenance window that it can use for regular updates. It is scheduled every Monday between 02:00 - 04:00 AM local time of the data center in which

the Managed Network Load Balancer service is deployed. During maintenance, a service interruption of up to 5 seconds may occur.

Additional update deployments may be possible and carried out outside the maintenance window, for example, when urgent security patches are needed.

> **Note:** IPv6 support has limitations for the Managed Network Load Balancer.

Managed Application Load Balancer

Next to the Managed Network Load Balancer, IONOS CLOUD also offers a Managed Application Load Balancer (ALB) that is balancing layer 7 HTTP or HTTPS-based network traffic. This service is available in all locations. While the Managed Network Load Balancer forwards the request according to the forwarding rule, the Managed Application Load Balancer will examine the content of the HTTP or HTTPS request header to determine where to route the request based on user-specified forwarding rules. Both types of Managed Load Balancers can be used in a stack for sequential traffic balancing.

Furthermore, the Managed Application Load Balancer can be provisioned as private and public load balancers (similar to the Managed Network Load Balancer).

The Managed Application Load Balancer supports HTTP and HTTPS protocols for forwarding rule configuration (incl. port specification). The configurations are managed in target groups which can contain multiple targets. An included SSL/ TLS certificate manager enables TLS offloading of the traffic. In addition, it can identify the server name, thus supporting multiple secure websites using one public IP but individual certificates.

The following types of routing are supported:

- Path-based routing
- hostname-based routing
- Routing support for query strings
- Header-based routing
- URL Redirection
- Static/Fixed response

When forwarding requests, the source IP can be preserved and forwarded in respective headers to the target so that the application can use information about the client that

actually made the request.

It offers options to specify health parameters to include or exclude nodes from the balancing configuration and manually remove listener targets from the load balancer (for example, when the node is in maintenance).

The number of Managed Application Load Balancers per contract is limited to five (5) ALBs. This limit can be adjusted by contacting the IONOS CLOUD Support Team.

The Managed Application Load Balancer supports the recording of Flow Logs.

The Managed Application Load Balancer operates in high-availability mode and service recovery is executed within seconds. The Managed Application Load Balancer will be regularly maintained by IONOS CLOUD and updated with the latest software versions and new features. IONOS CLOUD reserves a weekly maintenance window that it can use for regular updates. It is scheduled every Monday between 02:00 - 04:00 AM local time of the data center in which the Managed Application Load Balancer service is deployed. During maintenance, a service interruption of up to 5 seconds may occur.

Additional update deployments may be possible and carried out outside the maintenance window, for example, when urgent security patches are needed.

> **Note:** IPv6 support has limitations for the Managed Application Load Balancer.

Cloud DNS

IONOS CLOUD DNS allows customers to publish Domain Name System (DNS) zones for their domains and subdomains on public Name Servers.

Customers can manage their DNS zones and records via the Cloud DNS API and also create and manage Reverse DNS records for IPv4 and IPv6 addresses.

The IONOS CLOUD Name Server infrastructure is distributed across 14 points of presence (POPs) in Europe and the USA to ensure fast and reliable DNS resolution for customers in these locations.

Caching DNS

For the resolution of public domain names, IONOS CLOUD operates a redundant set consisting of two DNS servers at every data center location.

These DNS servers are operated as “caching” DNS servers or DNS resolvers, and are automatically assigned to the virtual customer entities via the DHCP IP address resolution.

Customer-specific internal domains cannot be resolved on caching DNS servers.

Reverse DNS

A standard reverse entry is assigned to all Public IPv4 Addresses, which are assigned to the virtual entities.

These entries follow the format ipAAA-BBB-CCC-DDD.pbiaas.com, whereby AAA-BBB-CCC-DDD corresponds to the IPv4 octets.

For statically assigned IPv4 and IPv6 addresses, the existing reverse entry can be adapted according to the customer requirements via Cloud DNS API.

Content Delivery Network

With the IONOS CLOUD Content Delivery Network (CDN), you can quickly deliver web content and applications to users with exceptional availability and performance. CDN offers a range of security features, including Layer 7 Distributed Denial of Service (DDoS) protection and a Web Application Firewall (WAF), making it an adaptable and secure solution for content delivery. CDN can be configured both via the DCD and CDN API.

IONOS CLOUD Content Delivery Network (CDN) is a network of servers located across the IONOS CLOUD global edge network to speed up the delivery of static and dynamic web content to users. CDN uses Anycast routing in IONOS CLOUD's global backbone network infrastructure with multiple highly available edge locations where content is distributed for reduced latency and high reliability of content loading on websites.

Currently, CDN hosts its edge locations in two European metro regions.

Network Security

Firewall

IONOS CLOUD allows the customer to use a software firewall within their infrastructure. For this purpose, the virtual network interface cards of a Virtual Server can be assigned a packet filter. The network traffic, which is aimed at the Virtual Server, is already filtered before the customer’s virtual machine.

Setting	Value
Throughput	Up to 6 Gbps

Network Security Groups

IONOS CLOUD enables customers to establish Network Security Groups (NSGs) that act as centralized policy managers for firewalls. These NSGs facilitate the filtering of network traffic both to and from virtual network resources within a specific Virtual Data Center (VDC).

Each NSG comprises security firewall rules that manage ingress (incoming) and egress (outgoing) network traffic. These rules apply to Network Interface Cards (NICs) and Virtual Machine (VM) resources that are configured as members of the NSG.

Network Security Groups are subject to the below limits

- Number of NSGs that can be created per VDC: 200
- Number of rules that can be created per NSG: 100
- Number of NSGs a VM can be a member of: 10
- Number of NSGs a NIC can be a member of: 10

DDoS Protect

DDoS Protect is a managed Distributed Denial of Service defense mechanism, which ensures that every customer resource hosted on IONOS CLOUD is secure and resilient against Layer 3 and Layer 4 DDoS attacks. This is facilitated by a filtering and scrubbing technology, which, if an attack is detected, filters the malicious DDoS traffic and lets through only the genuine traffic to its original destination. Hence, enabling applications and services of our customers to remain available under a DDoS attack.

Known attack vectors regularly evolve and new attack methods are added. IONOS CLOUD monitors this evolution and dedicates resources to adapt and enhance DDoS Protect as much as possible to capture and mitigate the threat.

The service is available in all our data centers and is a highly reliable, geo-distributed DDoS mitigation solution offering a robust 1.1 Tbit/s global filtering capacity, operating across 11 scrubbing blocks in 6 countries. The service is available in two packages:

DDoS Protect Basic: This package is enabled by default for all customers and does not require any configuration. It provides basic DDoS Protection for every resource on IONOS

CLOUD from common volumetric and protocol attacks and has the following features:

- DDoS traffic filtering - All suspicious traffic is redirected to the filtering platform where the DDoS traffic is filtered and the genuine traffic is allowed to the original destination.
- Always-On attack detection - The service is always on by default for all customers and does not require any added configuration or subscription.
- Automatic Containment - Each time an attack is identified the system automatically triggers the containment of the DDoS attack by activating the DDoS traffic and letting through only genuine traffic.
- Protection against common Layer 3 and 4 attacks - This service protects every resource on IONOS CLOUD from common volumetric and protocol attacks in the Network and Transport Layer such as UDP, SYN floods, etc.

DDoS Protect Advanced: This package offers everything that is part of the DDoS Protect Basic package plus advanced security measures and support.

- 24/7 DDoS Expert Support - Customers have 24/7 access to IONOS CLOUD DDoS expert support. The team is available to assist customers with their concerns regarding ongoing DDoS attacks or any related issues.
- Proactive Support - The IONOS CLOUD DDoS support team, equipped with alarms, will proactively respond to a DDoS attack directed towards a customer's resources and also notify the customer in such an event.
- On-demand IP specific DDoS filtering - If a customer suspects or anticipates a DDoS attack at any point in time, they can request to enable DDoS filtering for a specific IP or server owned by them. Once enabled, all traffic directed to that IP will be redirected to the IONOS CLOUD filtering platform where DDoS traffic will be filtered and genuine traffic will be passed to the original destination.
- On-demand Attack Diagnosis - At the customer's request, a detailed report of a DDoS attack is sent to the customer, explaining the attack and other relevant details.

NOTE: IONOS CLOUD sets forth Security as a Shared Responsibility between IONOS CLOUD and the customer. We at IONOS CLOUD strive at offering a state-of-the-art DDoS defense mechanism. Successful DDoS defense can only be achieved by a collective effort on all aspects including optimal use of firewalls and other settings in the customer environment.

Flow Logs

Flow log is a feature that allows you to capture data related to IPv4 and IPv6 network traffic flows. Flow logs can be enabled for any network interface of a virtual machine (VM) instance, the Managed Network Load Balancer, the Managed Application Load Balancer, and the public interfaces of the Managed Network Address Translation (NAT) Gateway.

Flow logs can help you with several tasks such as:

- Debugging connectivity and security issues
- Monitoring network throughput and performance
- Logging data to ensure that firewall rules are working as expected

The service can be configured for the direction of network traffic (ingress, egress, bi-directional) and action (accepted, rejected traffic packets, or any). Data is collected by the services and submitted in a compressed file to a customer's IONOS CLOUD Object Storage bucket, which can be specified by the customer at the time of flow log activation.

The service will not update existing files but will send new flow log records in a new compressed in an interval of 10 minutes.

Network Infrastructure Specifications

Core Network

IONOS CLOUD operates a high availability core network at each location for the redundant connection of the product platform. All services provided by IONOS CLOUD are connected to the Internet via this core network.

The core network consists exclusively of devices from brand manufacturers. The network connections are completed via an optical transmission network, which, by use of advanced technologies, can provide transmission capacities of several hundred gigabits per second. Connection to important Internet locations in Europe and America always guarantees the customer an optimal connection.

Data is not forwarded to third countries. At the customer’s explicit request, the customer can opt for support in a data center in a third country. In the interests of guaranteeing a suitable data protection level, this requires a separate agreement (within the meaning of article 44-50 DSGVO and §§ 78 ff. BDSG 2018).

Internal Network

IONOS CLOUD operates redundant networks at each location, offering connections up to 100 Gbps.

IONOS CLOUD uses high-speed networks based on InfiniBand technology both for connecting the central storage systems and for handling internal data connections between customer servers.

External Network

Depending on the location, different capacities for transmitting data to or from the Internet are available for operating the IONOS CLOUD service. Due to the direct connection between the data centers at the German locations, the upstream can be used across locations.

The total capacity of each respective location is described below:

Location	Connection	Location Redundancy	Metro Redundancy
Berlin (DE/TXL)	2 x 100 Gbps	2N	2N+1
Frankfurt am Main (DE/FRA)	2 x 100 Gbps	2N	2N+1
Frankfurt am Main (DE/FRA/2)	4 x 100 Gbps	2N	2N+1
Karlsruhe (DE/FKB)	2 x 100 Gbps	2N	2N+1
London (GB/LHR)	2 x 100 Gbps	2N	2N
Worcester (GB/BHX)	2 x 100 Gbps	2N	2N
Logroño (ES/VIT)	2 x 100 Gbps	2N	2N
Paris (FR/PAR)	2 x 100 Gbps	2N	2N+1
Las Vegas (US/LAS)	2 x 10 Gbps	2N	2N
Newark (US/EWR)	2 x 100 Gbps	2N	2N+1
Lenexa (US/MCI)	2 x 100 Gbps	2N	2N+1

IONOS CLOUD backbone AS-8560, to which IONOS CLOUD is redundantly connected, has a high quality edge capacity of 4000 Gbps with 3500 IPv4/IPv6 peering sessions, available in the following Internet and peering exchange points: AMS-IX, BW-IX,DE-CIX, NL-IX, ESPANIX, Equinix, FranceIX, KCIX, LINX.

Observability

Activity Log

For security audits, IONOS CLOUD offers a service called Activity Log. It traces and records activities of users when they log in, retrieve a resource, change, or delete resources. The trace records are limited to IaaS products. PaaS products will be expanded in near future iteratively.

Activity Log is available via RESTful API only. Users of role "Contract Owner" and "Administrators" are authorized by default to access the Activity Log API and are permitted to grant access to the Activity Log API for the accounts of type "User".

The data retention period for Activity Log is 35 days. After the retention period expires the data older than retention period gets purged. For longer persistence of Activity Log data, we recommend downloading the data and storing it on a different storage like IONOS CLOUD Object Storage.

Logging Service

Logging Service is designed to facilitate logging, monitoring, and analysis of logs from applications and infrastructure. It encompasses a diverse set of features aimed at enabling effective log monitoring and analysis, providing valuable insights into system behavior.

Using this service can enhance operational efficiency, bolster security measures, and elevate visibility into system performance and errors. To use the logging service, users must have the appropriate privileges, which can be configured through the DCD or Cloud API.

Logs can be sent from any location, including remote sources or systems outside the IONOS CLOUD Data centers.

The service provides a dedicated log server per logging pipeline instance to every customer or user. Each instance has a dedicated log server endpoint. While the Logging Service gets managed through a dedicated Logging API and a User Interface inside the DCD, it provides an endpoint to a Grafana dashboard to manage alerting and reporting.

The architecture of the Logging Service enables:

- Data Collection
- Data Aggregation
- Indexing and Storage
- Analysis and Reporting

The following log sources are currently supported:

- Kubernetes
- Docker
- Systemd
- HTTP (JSON REST API)
- Generic

Limitations

- Rate and Bandwidth limit per pipeline for log ingestion:
  - default HTTP rate limit: 50 requests per second
  - default TCP bandwidth: approximately 10,000 logs per second
- Maximum 5 pipelines per contract
- 10 log streams per pipeline.
- Only Linux Servers are supported (Windows Planned)

Available locations: Berlin, Frankfurt, Paris, Logroño, and London.

Monitoring Service

Monitoring Service is a cloud-based service that allows you to ingest, aggregate, and analyze data to enhance your understanding of your system's performance and behavior. The service collects data from various parts of your environment into a central system that is responsible for storage, aggregation, visualization, and automated responses when certain conditions are met.

The Monitoring Service can be managed only through the Monitoring Service API at the moment, Alerting and Reporting are available through a central Grafana instance, available to your contract per region.

Available locations: Berlin, Frankfurt, Paris, Logroño, and London.

- Berlin: `https://monitoring.de-txl.ionos.com/pipelines`
- Frankfurt: `https://monitoring.de-fra.ionos.com/pipelines`
- London: `https://monitoring.gb-lhr.ionos.com/pipelines`
- Paris: `https://monitoring.fr-par.ionos.com/pipelines`
- Logroño: `https://monitoring.es-vit.ionos.com/pipelines`

Contract limitations

Each contract can create 10 Monitoring Pipelines. If you require higher limitation boundaries, you can contact the IONOS CLOUD Support team to discuss your specific requirements and adjust the limits accordingly.

Sovereign Workplace

Nextcloud Workspace

Nextcloud Workspace is a managed, sovereign collaboration and productivity platform hosted securely in IONOS CLOUD data centers. It provides teams with a unified environment for secure communication, document management, and file sharing. By adhering to strict data protection standards and offering a sovereign infrastructure, the service caters to organizations requiring high compliance and data residency security. Delivered as a fully managed Software-as-a-Service (SaaS), IONOS CLOUD handles software maintenance, infrastructure scaling, security patching, and version updates.

Key Features

- Real-Time Collaboration: The platform provides a full-fledged online office suite (Collabora) for real-time collaborative document editing, including text processing, spreadsheets, and presentations.
- Cloud Storage: Each user receives 1 TB of shared cloud storage, featuring secure file sharing, file versioning, and optional end-to-end encryption.
- Unified Communication: Nextcloud Talk facilitates digital teamwork by natively combining team chat and professional video conferencing.
- Email Client: The platform features an integrated email client, allowing users to seamlessly connect and access their existing email provider through IMAP and SMTP protocols directly within the workspace.
- Private AI Assistant: An integrated AI assistant powered by the IONOS CLOUD AI Model Hub provides text generation, summarization, and translation services. AI inference runs within IONOS CLOUD infrastructure to contractually guarantee that customer data is never used to train models or exposed to third parties.
- Identity and Access Management (IAM): The platform supports enterprise standards such as LDAP, SAML, and OpenID Connect to securely integrate with existing customer

directories.
- Scalability: Currently, a single Nextcloud Workspace instance supports up to 200 users. We are actively working on scaling the service architecture to seamlessly accommodate the requirements of larger enterprise organizations in future releases.

Availability

Nextcloud Workspace is currently hosted in the IONOS CLOUD Berlin (TXL) data center and is available for global provisioning.

Security

Certificate Manager

The Certificate Manager streamlines the management of SSL/TLS certificates for your websites and applications, simplifying the processes of obtaining, installing, and renewing certificates. This ensures secure web traffic and the protection of sensitive data.

Key Features

- Simplified Management: Manage SSL/TLS certificates, ensuring efficient and consistent security across all your digital properties.
- Automated Workflows: Automate the issuance and renewal of certificates using ACME protocol-based domain validation, removing the hassle of manual uploads for certificates and keys.
- Seamless Integration: Effortlessly integrate with IONOS CLOUD services like Managed Application Load Balancer, CDN, and API Gateway to secure your workloads.

By providing these enhanced capabilities, the Certificate Manager helps secure your infrastructure while improving the efficiency of certificate management across cloud services.

SSH Key Support

Before provisioning, customers can inject the public part of an SSH key before provisioning using the IONOS CLOUD DCD or Cloud API (version 2.0 or higher) to create an SSH login for storage volumes based on a public IONOS CLOUD Linux image. This feature is not available

for snapshots, private images, Windows OS images, and MS SQL images provided by IONOS CLOUD.

In DCD, it is possible to store up to 100 public SSH keys for later re-use. It is possible to mark individual SSH keys as "default" which are applied to every provisioning automatically if previously mentioned requirements are fulfilled. Independent from the SSH key store, customers can also add SSH keys ad-hoc which they did not have previously stored. It is impossible to share access to an SSH key store across multiple users. Each user has access to their own SSH key store independent of their account role.

In Cloud API, it is only possible to add SSH keys ad-hoc. The SSH key store can't be accessed in this manner.

IONOS Private Cloud powered by VMware

IONOS Private Cloud powered by VMware brings VMware enterprise-class Software-Defined Data Center software to the IONOS CLOUD portfolio, enabling customers to run any application across vSphere-based private Cloud environments.

The Service Offering has the following components:

- Private Cloud consisting of:
  - VMware vSphere Enterprise Plus running on dedicated servers
  - VMware vCenter Server appliance
  - VMware NSX-T Standard to power networking for the Service Offering
  - VMware vSAN Standard aggregating host-based storage into a shared data store
- Self-service provisioning and scaling of Private Cloud on demand
- Access to the vSphere Client and NSX-T Manager, ability to configure and design the virtual data center that best fits ones use cases
- Maintenance, patching, and upgrades of the PC (the hardware, physical infrastructure, and VMware stack), performed by IONOS CLOUD

IONOS Private Cloud powered by VMware can be provisioned and scaled vertically on demand. A Private Cloud cluster includes a minimum of three hosts. Customers can add hosts, up to the provisioning maximum for their organization (maximum 24 hosts per cluster). Customers can select the location of the data center where their Private Cloud will be provisioned. Currently, the service is available in Berlin (Germany), Worcester (UK), and

Logroño (Spain). The Karlsruhe (AirPark) location will continue to operate already provisioned instances but does not allow the provisioning of new private cloud setups.

IONOS Private Cloud powered by VMware runs exclusively on high-performance, state-of-the-art hardware, and virtualization technology. vSAN, NSX-T, Intel Optane, and NVMe storage provide processing speeds that significantly optimize your workloads. The Backup Service can be used together with the Private Cloud. Additionally, we offer several other services that can be booked as an add-on.

Additional IP

For virtual network interface cards connected to the Internet, automatic assignment of a dynamic public IPv4 address through DHCP is standard. As an added advantage, IONOS CLOUD offers customers one static Public IPv4 Address included with their Private Cloud cluster. Further, customers can opt for additional static Public IPv4 Addresses at a nominal fee, facilitating assignment to Internet-connected virtual network interface cards as primary or supplementary IP addresses.

IPSec VPN

Connectivity to the vSphere management network of the Private Cloud is established through OpenVPN by default, allowing one terminal device connection at a time. A more comfortable alternative is provided by the service IPSec VPN. This service is integrated in the customer's Cloud Panel. This service leverages an automated provisioning of a Vyos backend gateway, enabling single or multiple IPSec connections to diverse endpoints. A singular service extends connectivity to multiple interconnected networks.

NFS datastore

Augmenting the local vSAN storage offered by the Private Cloud, the NFS datastore service allows customers to seamlessly incorporate network storage into their vSphere environment. This additional service permits scaling of storage from 2TB up to the standard provisioning maximum of 32TB (with an absolute maximum of 128TB). This flexibility empowers clients to scale their storage without the necessity of adding more hosts.

24/7 Enterprise Level Support

IONOS CLOUD offers its customers technical support by telephone and email. Experienced system administrators take the customers' calls and emails and address their concerns immediately. 24/7 Enterprise Level Support can be contacted by email or telephone 24 hours a day, seven days a week.

Contact details can be found in the **Data Center Designer dashboard**.

Public Cloud Locations

IONOS CLOUD uses data centers at different locations worldwide. The security concepts of data centers are always based on the highest industry standards.

Country / Location	Operator	Certifications
Germany / Berlin (DE/TXL)	IONOS CLOUD Group	BSI IT-Grundschutz, ISO/IEC 27001
Germany / Frankfurt am Main (DE/FRA)	Equinix GmbH	ISO/IEC 27001, ISO 9001, ISO 14001, ISO 22301, ISO 50001, SOC 1 Type II, SOC 2 Type II, PCI DSS
Germany / Frankfurt am Main (DE/FRA/2)	Digital Realty	ISO/IEC 27001, ISO 9001, ISO 14001, ISO 22301, ISO 50001, EN 50600,
Germany / Karlsruhe (DE/FKB)	TelemaxX Telekommunikation	GmbH BSI IT-Grundschutz, ISO/IEC 27001
United Kingdom / London (GB/LHR)	Virtus Data Centres Ltd.	ISO/IEC 27001, ISO 9001, ISO 14001, ISO 50001, ISO 20000, ISAE 3402 compliant, PCI-DSS compliant, BREEAM excellent
United Kingdom / Worcester (GB/BHX)	IONOS CLOUD Ltd.	Uptime Institute Tier IV, ISO/IEC 27001, ISO/IEC 50001, PCI-DSS
Spain / Logroño (ES/VIT)	Arsys Internet S.L.U.	AENOR ecommerce, ISO/IEC 27001, ISO 9001, ISO 50001, ENS Conformity, Uptime Tier III, and SAP Certified in Cloud and Infrastructure
France / Paris (FR/PAR)	Digital Realty	ISO14001, ISO22301, ISO27001, ISO50001, SOC2
USA / Las Vegas (US/LAS)	Switch, Inc.	ISO/IEC 27001, SSAE 16/ISAE 3402 (SOC1, SOC2), CICA SOC3, PCI-DSS
USA / Newark (US/EWR)	QTS Realty Trust, Inc.	ISO27001, SOC1, SOC2, PCI-DSS, HIPAA, FISMA
USA / Lenexa (US/MCI)	IONOS CLOUD Inc.	ISO/IEC 27001, ISO/IEC 50001

The Karlsruhe location is not available for the creation of new virtual data centers.

IONOS Private Cloud powered by VMware Locations

Country / Location	Operator	Certifications
Germany / Karlsruhe (AirPark)	IONOS CLOUD SE	ISO/IEC 27001, ISO/IEC 50001
Germany / Berlin	IONOS CLOUD Group	BSI IT-Grundschutz, ISO/IEC 27001
Spain / Logroño	Arsys Internet S.L.U.	ISO/IEC 27001, ISO/IEC 50001
France / Niederlauterbach	IONOS CLOUD Data Center SAS	ISO/IEC 27001, ISO/IEC 14001, ISO/IEC 50001
United Kingdom / Worcester	IONOS CLOUD Ltd.	Uptime Institute Tier IV, ISO/IEC 27001, ISO/IEC 50001, PCI-DSS

Since **5 July 2021**, it is no longer possible to create new Private Cloud servers at the Karlsruhe AirPark location.

Legacy Products

The following Cubes instances are not available for new deployments since October 1, 2024. They remain deployed and active for existing deployments:

Name	vCPU	RAM	DAS Storage
Cubes XS	1	1 GB	30 GB
Cubes S	1	2 GB	50 GB
Cubes M	2	4 GB	80 GB
Cubes L	4	8 GB	160 GB
Cubes XL	6	16 GB	320 GB
Cubes XXL	8	32 GB	640 GB
Cubes 3XL	12	48 GB	960 GB
Cubes 4XL	16	64 GB	1280 GB

3.3. Contact Information

Find phone numbers and email addresses for space.vars.ionos_cloud Support and Sales, organised by country and customer type.

Contact channels

To ensure we can assist you efficiently, we provide the following guidance on reaching the right team as quickly as possible.

IONOS Customers

For questions related to domains and hosting products managed via the **My IONOS Login** or **Cloud Panel portal**, including **VPS**, **Cloud Servers**, IONOS **Dedicated Server** and **Bare Metal Server**, please refer to the contact information at **ionos.co.uk/contact** or **ionos.de/kontakt**.

IONOS CLOUD Customers

For matters related to products or contracts managed via the Data Center Designer, **dcd.ionos.com**, or our VMware-based **Private Cloud product**, you can reach us via one of the following channels.

Please have your **Support PIN** and the associated email address at hand; for Private Cloud customers, please have the associated **Telephone PIN** ready.

3.4. Prices

Please check out the options below to view the price lists for IONOS products in your area.

- **IONOS CLOUD Inc.**
- **IONOS CLOUD Ltd.**
- **IONOS CLOUD (EUR)**
- **IONOS CLOUD SE (German)**

Thank you for choosing IONOS as your partner in navigating the digital world. We're here to support your success every step of the way.

IONOS CLOUD Inc.

The offer for the products of the IONOS CLOUD is for commercial use only. Prices are net prices and do not include any applicable taxes or duties, including but not limited to sales or value-added taxes.

As an IONOS CLOUD customer, you benefit from 24/7 phone and email technical support.

AI

AI Model Hub

LLMs	Price per million input tokens	Price per million output tokens	Price group
Llama 3.1 8B Instruct	$0.17	$0.17	PG A
Mistral Nemo Instruct	$0.17	$0.17	PG A
Mistral Small 24B Instruct	$0.11	$0.33	PG A
gpt-oss-120b	$0.17	$0.71	PG A
Llama 3.3 70B Instruct	$0.71	$0.71	PG A
Llama 3.1 405B Instruct	$1.93	$1.93	PG A

Code Models	Price per million input tokens	Price per million output tokens	Price group
Qwen3-Coder-Next (80B)	$0.17	$0.89	PG A
Code Llama 13B instruct HF	$0.50	$0.50	PG A

Please note: Code Llama 13B instruct HF is scheduled for deprecation on May 21, 2026.

Vision-Language Models	Price per million input tokens	Price per million output tokens	Price group
LightOnOCR 2	$0.17	$0.33	PG A

Text-to image models	Price per image	Price group
FLUX.1 [schnell]	$0.032	PG A

Embedding models	Price per million tokens	Price group
paraphrase-multilingual-mpnet-base-v2	$0.01	PG A
bge-large-en-v1.5	$0.015	PG A
bge-m3	$0.02	PG A
Qwen3-VL-Embedding-8B	$0.11	PG A
Qwen3-VL-Reranker-8B	$0.045	PG A

Document Collections (Legacy)

Vector database	Price per million tokens	Price group
Storage	$0.011 /30 days	PG A

Please note: The Document Collections feature is no longer available for new deployments.

Backup and Storage

Backup Service

Backup Storage Modes	Price per 30 days/GB	Price group
Cloud storage for orphaned backups	$0.082	PG F
Cloud storage for Windows or Mac WS backup	$0.082	PG F
Cloud storage for Windows or Linux server backup	$0.082	PG F
Cloud storage for virtual machine backup	$0.082	PG F
Backup storage on local machine	$0.082	PG F
Service fee for backup storage in IONOS CLOUD Object Storage	$0,017	PG F

Please note: Backup storage on Object Storage requires the subscription of the Direct Backup to Public Cloud feature pack.

Acronis Advanced Features	Price per 30 days	Price group
Advanced Security feature pack (per registered VM, server, workstation)	$2.00	PG A
Advanced Management feature pack (per registered VM, server, workstation)	$2.00	PG A
Direct Backup to Public Cloud feature pack (per registered VM, server, workstation)	$6.00	PG A

Block Storage

	Price per 30 days/GB	Price group
HDD	$0.0533	PG J
SSD Standard	$0.08	PG C2
SSD Premium	$0.17	PG C2
Snapshot	$0.0533	PG J

Free of charge Storage IOPS (Input/Output operations per second)

IONOS Cloud Object Storage

Storage	Price per 30 days	Price group
1 GB	$0.00487	PG C1

Data Storage Access	Price	Price group
Write	no charge	
Read	no charge	

API Call types	Price per 10.000 Calls	Price group
PUT/COPY/POST/LIST	no charge	
GET	no charge	
DELETE	no charge	

- 1 Gigabyte (GB) is equal to 1024 Megabytes (MB).
- Storage is billed with one-hour granularity.

Please note: Data transmission costs are not included and will be added according to the current pricing.

Network File Storage

Storage	Price per 30 days	Price group
1 TB	$169	PG B

Please note: the minimum provisioning size is 2 TB, up to a maximum of 42 TB.

Compute

Compute Engine

vCPU servers

CPU Type	Price per hour	Price group
vCPU	$0.013	PG C1

RAM	Price per hour	Price group
1 GB	$0.0022	PG C1

Dedicated Core servers

CPU Type	Price per hour	Price group
AMD Opteron*	$0.024	PG J
AMD EPYC Milan	$0.045	PG J
AMD EPYC Turin	$0.054	PG J
Intel® Core™ Haswell	$0.05	PG J
Intel® Skylake	$0.05	PG J
Intel® Ice Lake	$0.05	PG J
Intel® Sierra Forest	$0.056	PG J

RAM	Price per hour	Price group
1 GB	$0.0071	PG J

**AMD Opteron is not available for new deployments*

Cloud Savings Plans

The following hourly rates apply to resources covered by a 1 or 3-year Savings Plan.

Resource	Term	Price per hour	Price group
1 Dedicated Core	1 Year	$0.0370	PG C1
1 Dedicated Core	3 Years	$0.0261	PG D
1 GB RAM	1 Year	$0.0041	PG C1
1 GB RAM	3 Years	$0.0029	PG D

Please note:

1. Savings Plans cover all Dedicated Core CPU types except AMD Opteron.
2. Subscription to the Cloud Savings Plans is done through the DCD, subject to the acceptance of the Purchase Conditions.

Cloud GPU VM

Template	GPU Model	GPU Type	Number of GPUs	Dedicated CPUs	RAM (GiB)	NVMe Storage (GB)	Price per hour	Price group
H200-S	NVIDIA H200	H200 PCIe	1	15	267	1024	$3.261	PG C1
H200-M	NVIDIA H200	H200 PCIe	2	30	534	1536	$6.521	PG C1
H200-L	NVIDIA H200	H200 PCIe	4	60	1068	2048	$13.043	PG C1
H200-XL	NVIDIA H200	H200 PCIe	8	127	2136	4096	$26.086	PG C1

Microsoft licenses

Microsoft Windows Server (2016, 2019, 2022, 2025)

CPU type	Price per hour	Price group
Per vCPU	$0.0175	PG A
Per AMD Core	$0.0350	PG A
Per Intel® Core™	$0.0350	PG A

Microsoft SQL Server

Editions	Price per hour	Price group
Web Edition	$0.0194	PG A
Server Standard	$0.2639	PG A
Enterprise Edition	$1.19	PG A

The Microsoft SQL server license package includes the usage of 2 cores or 1 core pair. At least 2 license packages will be billed per server.

RedHat Enterprise Linux

RHEL on vCPU/Dedicated Core	Price per hour	Price group
RHEL Server Small Virtual Node (1-8 vCPUs/ Cores)	$0.0158 per vCPU or Core	PG A
RHEL Server Medium Virtual Node (9-127 vCPUs/ Cores)	$0.0118 per vCPU or Core	PG A
RHEL Server Large Virtual Node (128 or more vCPUs/ Cores)	$0.0105 per vCPU or Core	PG A

IONOS CLOUD Cubes

Available from October 1, 2024

Configuration	vCPU / RAM / Storage	Price per hour	Price group
Basic Cube XS	1 / 2 GB / 60 GB	$0.008	PG C1
Basic Cube S	2 / 4 GB / 120 GB	$0.014	PG C1
Basic Cube M	4 / 8 GB / 240 GB	$0.026	PG C1
Basic Cube L	8 / 16 GB / 480 GB	$0.048	PG C1
Basic Cube XL	16 / 32 GB / 960 GB	$0.091	PG C1
Memory Cube S	2 / 8 GB / 120 GB	$0.019	PG C1
Memory Cube M	4 / 16 GB/ 240 GB	$0.035	PG C1
Memory Cube L	8 / 32 GB / 480 GB	$0.067	PG C1
Memory Cube XL	16 / 64 GB / 960 GB	$0.125	PG C1

Available until September 30, 2024

Configuration	vCPU / RAM / Storage	Price per hour	Price group
Cubes XS	1 / 1 GB / 30 GB	$0.0031	PG C1
Cubes S	1 / 2 GB / 50 GB	$0.0061	PG C1
Cubes M	2 / 4 GB / 80 GB	$0.0107	PG C1
Cubes L	4 / 8 GB / 160 GB	$0.0229	PG C1
Cubes XL	6 / 16 GB / 320 GB	$0.0428	PG C1
Cubes XXL	8 / 32 GB / 640 GB	$0.0917	PG C1
Cubes 3XL	12 / 48 GB/ 960 GB	$0.1375	PG C1
Cubes 4XL	16 / 64 GB / 1280 GB	$0.1833	PG C1

Microsoft licenses

Microsoft Windows Server (2016, 2019, 2022, 2025)

License on Cubes	Price per hour	Price group
Per vCPU	$0.0079	PG A

RedHat Enterprise Linux

RHEL on Cubes	Price per hour	Price group
RHEL Server Small Virtual Node (1-8 vCPUs/ Cores)	$0.0158 per vCPU or Core	PG A
RHEL Server Medium Virtual Node (9-127 vCPUs/ Cores)	$0.0118 per vCPU or Core	PG A
RHEL Server Large Virtual Node (128 or more vCPUs/ Cores)	$0.0105 per vCPU or Core	PG A

Containers

Managed Kubernetes

New Managed Kubernetes Prices valid from April 1, 2025

Dedicated Core for Managed Kubernetes

CPU Type	Price per hour	Price group
Dedicated Core	$0.05	PG J

RAM	Price per hour	Price group
1 GB	$0.0071	PG J

Cloud Savings Plans

The following hourly rates apply to resources covered by a 1 or 3-year Savings Plan.

Resource	Term	Price per hour	Price group
1 Dedicated Core	1 Year	$0.0370	PG C1
1 Dedicated Core	3 Years	$0.0261	PG D
1 GB RAM	1 Year	$0.0041	PG C1
1 GB RAM	3 Years	$0.0029	PG D

Please note: Subscription to the Cloud Savings Plans is done through the DCD, subject to the acceptance of the Purchase Conditions.

vCPU for Managed Kubernetes

CPU Type	Price per hour	Price group
vCPU	$0.018	PG C1

RAM	Price per hour	Price group
1 GB	$0.0027	PG C1

Block Storage for Managed Kubernetes

	Price per 30 days/GB	Price group
HDD	$0.0533	PG J
SSD Premium	$0.17	PG C2

Private Container Registry

Storage	Price per 30 days	Price group
1 GB	$0.048	PG C1

Vulnerability Scanning	Price per 30 days	Price group
per 1 GB	$0.022	PG C1

Database as a Service

In-Memory DB

Core	Price	Price group
1	$0.066	PG H

RAM	Price per hour	Price group
1 GB	$0.008	PG H

Storage	Price per 30 days	Price group
1 GB SSD Premium	$0.17	PG C2

MariaDB

Core	Price per hour	Price group
1	$0.066	PG H

RAM	Price per hour	Price group
1 GB	$0.008	PG H

Storage	Price per 30 days	Price group
1 GB SSD Premium	$0.17	PG C2
1 GB Backup	$0.008	PG C1

MongoDB

Playground Edition

Instance Type	vCPU / RAM / Storage	Price per instance per hour	Price group
first instance	1 / 2 GB /50 GB	no charge	-
further instances	1/ 2 GB / 50 GB	$0,042	PG C1

Business Edition

Instance Type	vCPU / RAM / Storage	Price per Instance per hour	Price group
XS	1 / 2 GB / 50 GB	$0.042	PG C1
S	2 / 4 GB / 80 GB	$0.084	PG C1
M	4 / 8 GB / 160 GB	$0.168	PG C1
L	6 / 16 GB / 320 GB	$0.336	PG C1
XL	8 / 32 GB / 640 GB	$0.672	PG C1
XXL	16 / 64 GB / 960 GB	$1.356	PG C1
4XL	32 / 128 GB / 1280 GB	$2.724	PG C1
4XL_S	32 / 128 GB / 2000 GB	$2.964	PG C1

Enterprise Edition

Core	Price per hour	Price group
1	$0.066	PG C1

RAM	Price per hour	Price group
1 GB	$0.0388	PG C1

Storage	Price per 30 days	Price group
1 GB SSD Premium	$0.17	PG C2
1 GB SSD	$0.08	PG C2
1 GB HDD	$0.0533	PG C1

MongoDB Backup

Backup	Price per 30 days	Price group
1 GB Backup	$0.008	PG C1

PostgreSQL

Core	Price per hour	Price group
1	$0.072	PG H

RAM	Price per hour	Price group
1 GB	$0.0084	PG H

Storage	Price per 30 days	Price group
1 GB SSD Premium	$0.17	PG C2
1 GB SSD	$0.08	PG C2
1 GB HDD	$0.0533	PG H
1 GB Backup	$0.015 *[$0.008]**	PG C1

Data Analytics

Event Streams for Apache Kafka

Cluster size	Price per hour	Price group
XS	$0.35	PG H
S	$0.56	PG H
M	$0.74	PG H
L	$1.43	PG H
XL	$2.75	PG H

Network services

Data transfer

Traffic type	Volume per month	Price per 30 days/GB	Price group
Outgoing public traffic	up to 2 TB (≤ 2 TB)	no charge	PG J
Outgoing public traffic	next 8 TB (> 2 TB ≤ 10 TB)	$0.036	PG J
Outgoing public traffic	next 40 TB (> 10 TB ≤ 50 TB)	$0.030	PG J
Outgoing public traffic	next 100 TB (> 50 TB ≤ 150 TB)	$0.024	PG J
Outgoing public traffic	greater than 150 TB (> 150 TB)	$0.018	PG J
Outgoing public traffic	IONOS Public Cloud national and local	no charge	
Internal private traffic	-	no charge	
Incoming traffic	-	no charge	

IP Type	Price per 30 days	Price group
Reserved public IPv4	$6.00	PG J

Please note:

1. 1 TB equals 1,024 GB
2. Tiered pricing applies to all outgoing traffic from all IONOS Public Cloud services within the same contract, and the total cumulative amount is listed on the invoice.
3. Outgoing public traffic is classified in the following way:

3.1. IONOS Public Cloud local traffic is exchanged between services via public networks that terminate in the same IONOS Public Cloud data center;

3.2. IONOS Public Cloud national traffic is exchanged between services via public networks that terminate in a IONOS Public Cloud data center in the same country;

3.3. Outgoing data transfer between IONOS Public Cloud data centers in different countries, to IONOS Private Cloud, to IONOS Bare Metal Cloud as well as with the public internet, is charged as regular outgoing public traffic.

4. Internal private traffic is exchanged between services within the same IONOS Public Cloud data center via private networks.
5. IONOS CLOUD Object Storage:

Hosted across different locations, with contract-owned buckets and user-owned buckets options.

5.1 Traffic from a VM to Object Storage is subject to the same definition of local, national, and regular public traffic for cost invoicing:

5.1.1. Outgoing VM traffic to Object Storage in the same data center location is measured as local traffic;
5.1.2. Outgoing VM traffic to Object Storage in the same country, but to a different data center, is measured as national traffic;
5.1.3. Outgoing VM traffic to Object Storage in data centers in different countries is charged as regular public traffic.

5.2. Outgoing Object Storage traffic from **contract-owned buckets** is charged as regular public traffic, except for data transfer to services in the same data center or to IONOS Public Cloud data centers in the same country.

5.3. Outgoing Object Storage traffic from **user-owned buckets** is always charged as regular public traffic, and the definitions of local and national traffic do not apply, even

when this traffic is addressed to another IONOS Public Cloud service, regardless of its deployment location.

5.4. While inter-bucket data transfer is subject to charges, replication traffic both within the same region and across different regions is cost-free.

5.5. Outgoing Object Storage traffic to IONOS Bare Metal Cloud or IONOS Private Cloud is always charged as regular public traffic.

API Gateway

Please note: API Gateway is scheduled for deprecation and will reach its End-of-Life on December 31, 2025, after which it will no longer be available.

Requests	Price	Price group
Per 1 Million	$1.1	PG J

Please note: costs incurred for data transmission are not included, regular traffic charges apply.

Cloud DNS

	Price per 30 days	Price group
Public Zone (Primary)	$1.65/zone	PG J
Public Zone (Secondary)	$1.65/zone	PG J
Reverse DNS record (PTR)	no charge	

Content Delivery Network

	Price per 30 days	Price group
1 CDN	$109.00	PG A
Web Application Firewall	$1.00/routing rule	PG A

Managed Application Load Balancer

	Price per hour per rule	Price group
The first 5 forwarding rules	$0.024	PG J
Each additional forwarding rule	$0.012	PG J
Processing fee for incoming and outgoing data	no charge	

Please note: A minimum charge equivalent to one rule applies, even if none are configured.

Managed NAT Gateway

	Price per hour	Price group
Basis fee	$0.048	PG J
Processing fee for incoming and outgoing data	no charge	

Managed Network Load Balancer

	Price per hour per rule	Price group
The first 5 forwarding rules	$0.024	PG J
Each additional forwarding rule	$0.012	PG J
Processing fee for incoming and outgoing data	no charge	

Please note: A minimum charge equivalent to one rule applies, even if none are configured.

VPN Gateway

VPN Gateway type	Price per hour	Price group
Standard	$0.035	PG D
Standard (High Availability)	$0.07	PG D
Enhanced	$0.084	PG D
Enhanced (High Availability)	$0.165	PG D
Premium	$0.22	PG D
Premium (High Availability)	$0.44	PG D

Please note: price per provisioned Gateway. Costs incurred for data transmission are not included, regular traffic charges apply.

Observability

Logging Service

Component	Price	Price group
per Pipeline per hour	$0.02	PG A
per GB of storage per 30 days	$0.017	PG A

Monitoring Service

	Price	Price group
Per 1 Million samples	$0.17	PG J
per GB of storage per 30 days	$0.017	PG C1

Sovereign Workplace

Nextcloud Workspace

	Price per User / 30 days	Price group
Nextcloud Workspace (1 to 24 Users)	$8.44	PG G
Nextcloud Workspace (25 - 200 Users)	$7.13	PG G

Please note: Pay-As-You-Go pricing tiers are calculated independently for each individual Workspace. Each Workspace supports up to 200 users.

Cloud Savings Plans

The following rate applies to user seats covered by a 1-year Savings Plan. Savings Plans require a minimum of 25 users per plan and are applied at the contract level, meaning your plan's seats can be flexibly distributed across multiple Workspaces.

	Term	Price per User / 30 days	Price group
Nextcloud Workspace	1 Year	$6.24	PG E

v2026-05-13, prices valid from May 13, 2026

IONOS CLOUD Inc., 200 Continental Drive Suite 401, Newark, DE 19713

IONOS CLOUD Ltd.

The offer for IONOS CLOUD products is for commercial use only. Prices are net prices and do not include any applicable taxes or duties, including but not limited to sales or value-added taxes.

As an IONOS CLOUD customer, you benefit from 24/7 phone and email technical support.

AI

AI Model Hub

Large Language Models	Price per million input tokens	Price per million output tokens	Price group
Llama 3.1 8B Instruct	£0.14	£0.14	PG A
Mistral Nemo Instruct	£0.14	£0.14	PG A
Mistral Small 24B Instruct	£0.085	£0.26	PG A
gpt-oss-120b	£0.14	£0.55	PG A
Llama 3.3 70B Instruct	£0.55	£0.55	PG A
Llama 3.1 405B Instruct	£1.58	£1.58	PG A

Code Models	Price per million input tokens	Price per million output tokens	Price group
Qwen3-Coder-Next (80B)	£0.13	£0.68	PG A
Code Llama 13B instruct HF	£0.41	£0.41	PG A

Please note: Code Llama 13B instruct HF is scheduled for deprecation on May 21, 2026.

Vision-Language Models	Price per million input tokens	Price per million output tokens	Price group
LightOnOCR 2	£0.13	£0.26	PG A

Text-to image models	Price per image	Price group
FLUX.1 [schnell]	£0.0256	PG A

Embedding models	Price per million tokens	Price group
paraphrase-multilingual-mpnet-base-v2	£0.01	PG A
bge-large-en-v1.5	£0.015	PG A
bge-m3	£0.02	PG A
Qwen3-VL-Embedding-8B	£0.09	PG A
Qwen3-VL-Reranker-8B	£0.034	PG A

Document Collections (Legacy)

Vector database	Price per million tokens	Price group
Storage	£0.009/30 days	PG A

Please note: The Document Collections feature is no longer available for new deployments.

Backup and Storage

Backup Service

Backup Storage Modes	Price per 30 days/GB	Price group
Cloud storage for orphaned backups	£0.07	PG F
Cloud storage for Windows or Mac WS backup	£0.07	PG F
Cloud storage for Windows or Linux server backup	£0.07	PG F
Cloud storage for virtual machine backup	£0.07	PG F
Backup storage on local machine	£0.07	PG F
Service fee for backup storage in IONOS CLOUD Object Storage	£0,014	PG F

Please note: Backup storage on Object Storage requires the subscription of the Direct Backup to Public Cloud feature pack.

Acronis Advanced Features	Price per 30 days	Price group
Advanced Security feature pack (per registered VM, server, workstation)	£1.50	PG A
Advanced Management feature pack (per registered VM, server, workstation)	£1.50	PG A
Direct Backup to Public Cloud feature pack (per registered VM, server, workstation)	£5.00	PG A

Block Storage

	Price per 30 days/GB	Price group
HDD	£0.04	PG J
SSD Standard	£0,063	PG C2
SSD Premium	£0,14	PG C2
Snapshot	£0.04	PG J

Free of charge Storage IOPS (Input/Output operations per second)

IONOS Cloud Object Storage

Storage	Price per 30 days/GB	Price group
1 GB	£0,0063	PG C1

API Call types	Price per 10.000 Calls	Price group
PUT/COPY/POST/LIST	no charge	
GET	no charge	
DELETE	no charge	

Data Storage Access	Price per GB	Price group
Write	no charge	
Read	no charge	

- 1 Gigabyte (GB) is equal to 1,024 Megabytes (MB).
- Storage is billed with one-hour granularity.

Please note: Data transmission costs are not included and will be added according to the current pricing.

Network File Storage

Storage	Price per 30 days	Price group
1 TB	£139	PG B

Please note: the minimum provisioning size is 2 TB, up to a maximum of 42 TB.

Compute

Compute Engine

vCPU servers

CPU type	Price per hour	Price group
vCPU	£0.011	PG C1

RAM	Price per hour	Price group
1 GB	£0.0018	PG C1

Dedicated Core servers

CPU type	Price per hour	Price group
AMD Opteron*	£0.02	PG J
AMD EPYC Milan	£0.036	PG J
AMD EPYC Turin	£0.043	PG J
Intel® Core™ Haswell	£0.04	PG J
Intel® Core™ Skylake	£0.04	PG J
Intel® Core™ Ice Lake	£0.04	PG J
Intel® Core™ Sierra Forest	£0.047	PG J

RAM	Price per hour	Price group
1 GB	£0.0045	PG J

**AMD Opteron is not available for new deployments*

Cloud Savings Plans

The following hourly rates apply to resources covered by a 1 or 3-year Savings Plan.

Resource	Term	Price per hour	Price group
1 Dedicated Core	1 Year	£0.0290	PG C1
1 Dedicated Core	3 Years	£0.0204	PG D
1 GB RAM	1 Year	£0.0032	PG C1
1 GB RAM	3 Years	£0.0023	PG D

Please note:

1. Savings Plans cover all Dedicated Core CPU types except AMD Opteron.
2. Subscription to the Cloud Savings Plans is done through the DCD, subject to the acceptance of the Purchase Conditions.

Cloud GPU VM

Template	GPU Model	GPU Type	Number of GPUs	Dedicated CPUs	RAM (GiB)	NVMe Storage (GB)	Price per hour	Price group
H200-S	NVIDIA H200	H200 PCIe	1	15	267	1024	£2.554	PG C1
H200-M	NVIDIA H200	H200 PCIe	2	30	534	1536	£5.109	PG C1
H200-L	NVIDIA H200	H200 PCIe	4	60	1068	2048	£10.218	PG C1
H200-XL	NVIDIA H200	H200 PCIe	8	127	2136	4096	£20.436	PG C1

Microsoft licences

Microsoft Windows Server (2016, 2019, 2022, 2025)

CPU type	Price per hour	Price group
Per vCPU	£0.0145	PG A
Per AMD Core	£0.0292	PG A
Per Intel® Core™	£0.0292	PG A

Microsoft SQL Server

Editions	Price per hour	Price group
Web Edition	£0.0167	PG A
Server Standard	£0.2222	PG A
Enterprise Edition	£0.99	PG A

The Microsoft SQL server licence package includes the usage of 2 cores or 1 core pair. At least 2 licence packages will be billed per server.

RedHat Enterprise Linux

RHEL on vCPU/Dedicated Core	Price per hour	Price group
RHEL Server Small Virtual Node (1-8 vCPUs/ Cores)	£0.0129 per vCPU or Core	PG A
RHEL Server Medium Virtual Node (9-127 vCPUs/ Cores)	£0.0096 per vCPU or Core	PG A
RHEL Server Large Virtual Node (128 or more vCPUs/ Cores)	£0.0085 per vCPU or Core	PG A

IONOS CLOUD Cubes

Available from 1 October 2024

Configuration	vCPU / RAM / Storage	Price per hour	Price group
Basic Cube XS	1 / 2 GB / 60 GB	£0.007	PG C1
Basic Cube S	2 / 4 GB / 120 GB	£0.013	PG C1
Basic Cube M	4 / 8 GB / 240 GB	£0.024	PG C1
Basic Cube L	8 / 16 GB / 480 GB	£0.044	PG C1
Basic Cube XL	16 / 32 GB / 960 GB	£0.083	PG C1
Memory Cube S	2 / 8 GB / 120 GB	£0.017	PG C1
Memory Cube M	4 / 16 GB/ 240 GB	£0.032	PG C1
Memory Cube L	8 / 32 GB / 480 GB	£0.061	PG C1
Memory Cube XL	16 / 64 GB / 960 GB	£0.114	PG C1

Available until 30 September 2024

Configuration	vCPU / RAM / Storage	Price per hour	Price group
Cubes XS	1 / 1 GB / 30 GB	£0.0025	PG C1
Cubes S	1 / 2 GB / 50 GB	£0.005	PG C1
Cubes M	2 / 4 GB / 80 GB	£0.0088	PG C1
Cubes L	4 / 8 GB / 160 GB	£0.0188	PG C1
Cubes XL	6 / 16 GB / 320 GB	£0.035	PG C1
Cubes XXL	8 / 32 GB / 640 GB	£0.075	PG C1
Cubes 3XL	12 / 48 GB/ 960 GB	£0.1125	PG C1
Cubes 4XL	16 / 64 GB / 1280 GB	£0.15	PG C1

Microsoft licences

Microsoft Windows Server (2016, 2019, 2022, 2025)

Licence on Cubes	Price per hour	Price group
Per vCPU	£0.0066	PG A

RedHat Enterprise Linux

RHEL on Cubes	Price per hour	Price group
RHEL Server Small Virtual Node (1-8 vCPUs/ Cores)	£0.0129 per vCPU or Core	PG A
RHEL Server Medium Virtual Node (9-127 vCPUs/ Cores)	£0.0096 per vCPU or Core	PG A
RHEL Server Large Virtual Node (128 or more vCPUs/ Cores)	£0.0085 per vCPU or Core	PG A

Containers

Managed Kubernetes

New Managed Kubernetes Prices valid from April 1, 2025

Dedicated Core for Managed Kubernetes

CPU Type	Price per hour	Price group
Dedicated Core	£0.04	PG J

RAM	Price per hour	Price group
1 GB	£0.0045	PG J

Cloud Savings Plans

The following hourly rates apply to resources covered by a 1 or 3-year Savings Plan.

Resource	Term	Price per hour	Price group
1 Dedicated Core	1 Year	£0.0290	PG C1
1 Dedicated Core	3 Years	£0.0204	PG D
1 GB RAM	1 Year	£0.0032	PG C1
1 GB RAM	3 Years	£0.0023	PG D

Please note: Subscription to the Cloud Savings Plans is done through the DCD, subject to the acceptance of the Purchase Conditions.

vCPU for Managed Kubernetes

CPU Type	Price per hour	Price group
vCPU	£0.016	PG C1

RAM	Price per hour	Price group
1 GB	£0.0023	PG C1

Block Storage for Managed Kubernetes

	Price per 30 days/GB	Price group
HDD	£0.04	PG J
SSD Premium	£0.14	PG C2

Private Container Registry

Storage	Price per 30 days	Price group
1 GB	£0.04	PG C1
Vulnerability Scanning	Price per 30 days	Price group
-	-	-
per 1 GB	£0.018	PG C1

Database as a Service

In-Memory DB

Core	Price per hour	Price group
1	£0.054	PG H

RAM	Price per hour	Price group
1 GB	£0.006	PG H

Storage	Price per 30 days	Price group
1 GB SSD Premium	£0.14	PG C2

MariaDB

Core	Price per hour	Price group
1	£0.054	PG H

RAM	Price per hour	Price group
1 GB	£0.006	PG H

Storage	Price per 30 days	Price group
1 GB SSD Premium	£0.14	PG C2
1 GB Backup	£0.0063	PG C1

MongoDB

Playground Edition

Instance Type	vCPU / RAM / Storage	Price per hour	Price group
first instance	1 / 2 GB /50 GB	no charge	-
further instances	1/ 2 GB / 50 GB	£0.035	PG C1

Business Edition

Instance Type	vCPU / RAM / Storage	Price per hour	Price group
XS	1 / 2 GB / 50 GB	£0.035	PG C1
S	2 / 4 GB / 80 GB	£0.070	PG C1
M	4 / 8 GB / 160 GB	£0.140	PG C1
L	6 / 16 GB / 320 GB	£0.280	PG C1
XL	8 / 32 GB / 640 GB	£0.560	PG C1
XXL	16 / 64 GB / 960 GB	£1.130	PG C1
4XL	32 / 128 GB / 1280 GB	£2.270	PG C1
4XL_S	32 / 128 GB / 2000 GB	£2.470	PG C1

Enterprise Edition

Core	Price per hour	Price group
1	£0.054	PG C1

RAM	Price per hour	Price group
1 GB	£0.0318	PG C1

Storage	Price per 30 days	Price group
1 GB SSD Premium	£0.14	PG C2
1 GB SSD	0.063	PG C2
1 GB HDD	£0.04	PG C1

MongoDB Backup

Backup	Price per 30 days	Price group
1 GB Backup	£0.0063	PG C1

PostgreSQL

Core	Price per hour	Price group
1	£0.06	PG H

RAM	Price per hour	Price group
1 GB	£0.007	PG H

Storage	Price per 30 days	Price group
1 GB SSD Premium	£0.14	PG C2
1 GB SSD	£0.063	PG C2
1 GB HDD	£0.04	PG H
1 GB Backup	£0.0063	PG C1

Data Analytics

Event Streams for Apache Kafka

Cluster size	Price per hour	Price group
XS	£0.29	PG H
S	£0.46	PG H
M	£0.60	PG H
L	£1.2	PG H
XL	£2.3	PG H

Network services

Data transfer

Traffic type	Volume per month	Price per 30 days/GB	Price group
Outgoing public traffic	up to 2 TB (≤ 2 TB)	no charge	PG J
Outgoing public traffic	next 8 TB (> 2 TB ≤ 10 TB)	£0.030	PG J
Outgoing public traffic	next 40 TB(> 10 TB ≤ 50 TB)	£0.025	PG J
Outgoing public traffic	next 100 TB(> 50 TB ≤ 150 TB)	£0.020	PG J
Outgoing public traffic	greater than 150 TB (> 150 TB)	£0.015	PG J
Outgoing public traffic	IONOS Public Cloud national and local	no charge	
Internal private traffic	-	no charge	
Incoming traffic	-	no charge	

IP Type	Price per 30 days	Price group
Reserved public IPv4	£5.00	PG J

Please note:

1. 1 TB equals 1,024 GB
2. Tiered pricing applies to all outgoing traffic from all IONOS Public Cloud services within the same contract, and the total cumulative amount is listed on the invoice.
3. Outgoing public traffic is classified in the following way:

3.1. IONOS Public Cloud local traffic is exchanged between services via public networks that terminate in the same IONOS Public Cloud data centre;

3.2. IONOS Public Cloud national traffic is exchanged between services via public networks that terminate in a IONOS Public Cloud data centre in the same country;

3.3. Outgoing data transfer between IONOS Public Cloud data centres in different countries, to IONOS Private Cloud, to IONOS Bare Metal Cloud as well as with the public internet, is charged as regular outgoing public traffic.

4. Internal private traffic is exchanged between services within the same IONOS Public Cloud data centre via private networks.
5. IONOS CLOUD Object Storage:

Hosted across different locations, with contract-owned buckets and user-owned buckets options.

5.1 Traffic from a VM to Object Storage is subject to the same definition of local, national, and regular public traffic for cost invoicing:

5.1.1. Outgoing VM traffic to Object Storage in the same data centre location is measured as local traffic;
5.1.2. Outgoing VM traffic to Object Storage in the same country, but to a different data centre, is measured as national traffic;
5.1.3. Outgoing VM traffic to Object Storage in data centres in different countries is charged as regular public traffic.

5.2. Outgoing Object Storage traffic from **contract-owned buckets** is charged as regular public traffic, except for data transfer to services in the same data centre or to IONOS Public Cloud data centres in the same country.

5.3. Outgoing Object Storage traffic from **user-owned buckets** is always charged as regular public traffic, and the definitions of local and national traffic do not apply, even

when this traffic is addressed to another IONOS Public Cloud service, regardless of its deployment location.

5.4. While inter-bucket data transfer is subject to charges, replication traffic both within the same region and across different regions is cost-free.

5.5. Outgoing Object Storage traffic to IONOS Bare Metal Cloud or IONOS Private Cloud is always charged as regular public traffic.

Cloud DNS

	Price per 30 days	Price group
Public Zone (Primary)	£1.35/zone	PG J
Public Zone (Secondary)	£1.35/zone	PG J
Reverse DNS record (PTR)	no charge	

Content Delivery Network

	Price per 30 days	Price group
1 CDN	£89.00	PG A
Web Application Firewall	£0.80/routing rule	PG A

Managed Application Load Balancer

	Price per hour per rule	Price group
The first 5 forwarding rules	£0.02	PG J
Each additional forwarding rule	£0.01	PG J
Processing fee for incoming and outgoing data	no charge	

Please note: A minimum charge equivalent to one rule applies, even if none are configured.

Managed NAT Gateway

	Price per hour	Price group
Basis fee	£0.04	PG J
Processing fee for incoming and outgoing data	no charge	

Managed Network Load Balancer

	Price per hour per rule	Price group
The first 5 forwarding rules	£0.02	PG J
Each additional forwarding rule	£0.01	PG J
Processing fee for incoming and outgoing data	no charge	

Please note: A minimum charge equivalent to one rule applies, even if none are configured.

VPN Gateway

VPN Gateway type	Price per hour	Price group
Standard	£0.03	PG D
Standard (High Availability)	£0.06	PG D
Enhanced	£0.07	PG D
Enhanced (High Availability)	£0.14	PG D
Premium	£0.18	PG D
Premium (High Availability)	£0.36	PG D

Please note: price per provisioned Gateway. Costs incurred for data transmission are not included, regular traffic charges apply.

Observability

Logging Service

Component	Price	Price group
per Pipeline per hour	£0.016	PG A
per GB of storage per 30 days	£0.014	PG A

Monitoring Service

	Price	Price group
Per 1 Million samples	£0.14	PG J
per GB of storage per 30 days	£0.014	PG C1

Sovereign Workplace

Nextcloud Workspace

	Price per User / 30 days	Price group
Nextcloud Workspace (1 to 24 Users)	£6.43	PG G
Nextcloud Workspace (25+ Users)	£5.43	PG G

Please note: Pay-As-You-Go pricing tiers are calculated independently for each individual Workspace. Each Workspace supports up to 200 users.

Cloud Savings Plans

The following rate applies to user seats covered by a 1-year Savings Plan. Savings Plans require a minimum of 25 users per plan and are applied at the contract level, meaning your plan's seats can be flexibly distributed across multiple Workspaces.

	Term	Price per User / 30 days	Price group
Nextcloud Workspace	1 Year	£4.75	PG E

Private Cloud powered by VMware

Basic bundle (Cluster consisting of 3 hosts, incl. VMware licences)	Price per month	Price group
Silver-48	£3,395.00	PG 1p
Silver-96	£4,600.00	PG 1p
Gold-192 v2	£5,380.00	PG 1p
Gold-512	£7,385.00	PG 1p
M-512	£5,400.00	PG 1p
L-768	£8,000.00	PG 1p
XL-1024	£10,900.00	PG 1p

Additional host (incl. VMware licences)	Price per month	Price group
Silver-48	£740.00	PG 1p
Silver-96	£915.00	PG 1p
Gold-192 v2	£1,350.00	PG 1p
Gold-512	£2,220.00	PG 1p
M-512	£1,400.00	PG 1p
L-768	£2,300.00	PG 1p
XL-1024	£3,100.00	PG 1p

Traffic type	Price	Price group
Outgoing	no charge	PG 1p
Incoming	no charge	PG 1p
National	no charge	PG 1p

Backup	Price	Price group
Cloud Backup	£/GB 0.07	PG 1p

Additional IP address(es)	Price per month	Price group
1 Additional IP address	£5.00	PG 1p
Subnet/28 (16)	£24.00	PG 1p
Subnet/27 (32)	£32.00	PG 1p

Network services	Price per month	Price group
IPSEC VPN	£50.00	PG 1p

Additional storage	Price per month	Price group
NFS Datastore (min. 2 TB)	£72.00/TB	PG 1p

VCDA Protection	Price per month	Price group
VCDA Protection	£50.00/VM	PG 1p

Please note: Only VM's that are protected will be invoiced, the migration feature can be used totally free.

Change hardware model	Price (One-Time Fee)
Change hardware model	£300.00

Please note: Outgoing traffic to IONOS Bare Metal Cloud and Private Cloud powered by VMware is recorded as regular public traffic.

v2026-05-13, prices valid from 13/05/2026

IONOS CLOUD Limited, Discovery House, 154 Southgate Street, Gloucester, GL1 2EX, United Kingdom Registered in England and Wales, Company number 3953678 - VAT No GB 752539027

3.4.3. IONOS CLOUD (EUR)

IONOS CLOUD Ltd.

The offer for IONOS CLOUD products is for commercial use only. Prices are net prices and do not include any applicable taxes or duties, including but not limited to sales or value-added taxes.

As an IONOS CLOUD customer, you benefit from 24/7 phone and email technical support.

AI

AI Model Hub

Large Language Models	Price per million input tokens	Price per million output tokens	Price group
Llama 3.1 8B Instruct	£0.14	£0.14	PG A
Mistral Nemo Instruct	£0.14	£0.14	PG A
Mistral Small 24B Instruct	£0.085	£0.26	PG A
gpt-oss-120b	£0.14	£0.55	PG A
Llama 3.3 70B Instruct	£0.55	£0.55	PG A
Llama 3.1 405B Instruct	£1.58	£1.58	PG A

Code Models	Price per million input tokens	Price per million output tokens	Price group
Qwen3-Coder-Next (80B)	£0.13	£0.68	PG A
Code Llama 13B instruct HF	£0.41	£0.41	PG A

Please note: Code Llama 13B instruct HF is scheduled for deprecation on May 21, 2026.

Vision-Language Models	Price per million input tokens	Price per million output tokens	Price group
LightOnOCR 2	£0.13	£0.26	PG A

Text-to image models	Price per image	Price group
FLUX.1 [schnell]	£0.0256	PG A

Embedding models	Price per million tokens	Price group
paraphrase-multilingual-mpnet-base-v2	£0.01	PG A
bge-large-en-v1.5	£0.015	PG A
bge-m3	£0.02	PG A
Qwen3-VL-Embedding-8B	£0.09	PG A
Qwen3-VL-Reranker-8B	£0.034	PG A

Document Collections (Legacy)

Vector database	Price per million tokens	Price group
Storage	£0.009/30 days	PG A

Please note: The Document Collections feature is no longer available for new deployments.

Backup and Storage

Backup Service

Backup Storage Modes	Price per 30 days/GB	Price group
Cloud storage for orphaned backups	£0.07	PG F
Cloud storage for Windows or Mac WS backup	£0.07	PG F
Cloud storage for Windows or Linux server backup	£0.07	PG F
Cloud storage for virtual machine backup	£0.07	PG F
Backup storage on local machine	£0.07	PG F
Service fee for backup storage in IONOS CLOUD Object Storage	£0,014	PG F

Please note: Backup storage on Object Storage requires the subscription of the Direct Backup to Public Cloud feature pack.

Acronis Advanced Features	Price per 30 days	Price group
Advanced Security feature pack (per registered VM, server, workstation)	£1.50	PG A
Advanced Management feature pack (per registered VM, server, workstation)	£1.50	PG A
Direct Backup to Public Cloud feature pack (per registered VM, server, workstation)	£5.00	PG A

Block Storage

	Price per 30 days/GB	Price group
HDD	£0.04	PG J
SSD Standard	£0,063	PG C2
SSD Premium	£0,14	PG C2
Snapshot	£0.04	PG J

Free of charge Storage IOPS (Input/Output operations per second)

IONOS Cloud Object Storage

Storage	Price per 30 days/GB	Price group
1 GB	£0,0063	PG C1

API Call types	Price per 10.000 Calls	Price group
PUT/COPY/POST/LIST	no charge	
GET	no charge	
DELETE	no charge	

Data Storage Access	Price per GB	Price group
Write	no charge	
Read	no charge	

- 1 Gigabyte (GB) is equal to 1,024 Megabytes (MB).
- Storage is billed with one-hour granularity.

Please note: Data transmission costs are not included and will be added according to the current pricing.

Network File Storage

Storage	Price per 30 days	Price group
1 TB	£139	PG B

Please note: the minimum provisioning size is 2 TB, up to a maximum of 42 TB.

Compute

Compute Engine

vCPU servers

CPU type	Price per hour	Price group
vCPU	£0.011	PG C1

RAM	Price per hour	Price group
1 GB	£0.0018	PG C1

Dedicated Core servers

CPU type	Price per hour	Price group
AMD Opteron*	£0.02	PG J
AMD EPYC Milan	£0.036	PG J
AMD EPYC Turin	£0.043	PG J
Intel® Core™ Haswell	£0.04	PG J
Intel® Core™ Skylake	£0.04	PG J
Intel® Core™ Ice Lake	£0.04	PG J
Intel® Core™ Sierra Forest	£0.047	PG J

RAM	Price per hour	Price group
1 GB	£0.0045	PG J

**AMD Opteron is not available for new deployments*

Cloud Savings Plans

The following hourly rates apply to resources covered by a 1 or 3-year Savings Plan.

Resource	Term	Price per hour	Price group
1 Dedicated Core	1 Year	£0.0290	PG C1
1 Dedicated Core	3 Years	£0.0204	PG D
1 GB RAM	1 Year	£0.0032	PG C1
1 GB RAM	3 Years	£0.0023	PG D

Please note:

1. Savings Plans cover all Dedicated Core CPU types except AMD Opteron.
2. Subscription to the Cloud Savings Plans is done through the DCD, subject to the acceptance of the Purchase Conditions.

Cloud GPU VM

Template	GPU Model	GPU Type	Number of GPUs	Dedicated CPUs	RAM (GiB)	NVMe Storage (GB)	Price per hour	Price group
H200-S	NVIDIA H200	H200 PCIe	1	15	267	1024	£2.554	PG C1
H200-M	NVIDIA H200	H200 PCIe	2	30	534	1536	£5.109	PG C1
H200-L	NVIDIA H200	H200 PCIe	4	60	1068	2048	£10.218	PG C1
H200-XL	NVIDIA H200	H200 PCIe	8	127	2136	4096	£20.436	PG C1

Microsoft licences

Microsoft Windows Server (2016, 2019, 2022, 2025)

CPU type	Price per hour	Price group
Per vCPU	£0.0145	PG A
Per AMD Core	£0.0292	PG A
Per Intel® Core™	£0.0292	PG A

Microsoft SQL Server

Editions	Price per hour	Price group
Web Edition	£0.0167	PG A
Server Standard	£0.2222	PG A
Enterprise Edition	£0.99	PG A

The Microsoft SQL server licence package includes the usage of 2 cores or 1 core pair. At least 2 licence packages will be billed per server.

RedHat Enterprise Linux

RHEL on vCPU/Dedicated Core	Price per hour	Price group
RHEL Server Small Virtual Node (1-8 vCPUs/ Cores)	£0.0129 per vCPU or Core	PG A
RHEL Server Medium Virtual Node (9-127 vCPUs/ Cores)	£0.0096 per vCPU or Core	PG A
RHEL Server Large Virtual Node (128 or more vCPUs/ Cores)	£0.0085 per vCPU or Core	PG A

IONOS CLOUD Cubes

Available from 1 October 2024

Configuration	vCPU / RAM / Storage	Price per hour	Price group
Basic Cube XS	1 / 2 GB / 60 GB	£0.007	PG C1
Basic Cube S	2 / 4 GB / 120 GB	£0.013	PG C1
Basic Cube M	4 / 8 GB / 240 GB	£0.024	PG C1
Basic Cube L	8 / 16 GB / 480 GB	£0.044	PG C1
Basic Cube XL	16 / 32 GB / 960 GB	£0.083	PG C1
Memory Cube S	2 / 8 GB / 120 GB	£0.017	PG C1
Memory Cube M	4 / 16 GB/ 240 GB	£0.032	PG C1
Memory Cube L	8 / 32 GB / 480 GB	£0.061	PG C1
Memory Cube XL	16 / 64 GB / 960 GB	£0.114	PG C1

Available until 30 September 2024

Configuration	vCPU / RAM / Storage	Price per hour	Price group
Cubes XS	1 / 1 GB / 30 GB	£0.0025	PG C1
Cubes S	1 / 2 GB / 50 GB	£0.005	PG C1
Cubes M	2 / 4 GB / 80 GB	£0.0088	PG C1
Cubes L	4 / 8 GB / 160 GB	£0.0188	PG C1
Cubes XL	6 / 16 GB / 320 GB	£0.035	PG C1
Cubes XXL	8 / 32 GB / 640 GB	£0.075	PG C1
Cubes 3XL	12 / 48 GB/ 960 GB	£0.1125	PG C1
Cubes 4XL	16 / 64 GB / 1280 GB	£0.15	PG C1

Microsoft licences

Microsoft Windows Server (2016, 2019, 2022, 2025)

Licence on Cubes	Price per hour	Price group
Per vCPU	£0.0066	PG A

RedHat Enterprise Linux

RHEL on Cubes	Price per hour	Price group
RHEL Server Small Virtual Node (1-8 vCPUs/ Cores)	£0.0129 per vCPU or Core	PG A
RHEL Server Medium Virtual Node (9-127 vCPUs/ Cores)	£0.0096 per vCPU or Core	PG A
RHEL Server Large Virtual Node (128 or more vCPUs/ Cores)	£0.0085 per vCPU or Core	PG A

Containers

Managed Kubernetes

New Managed Kubernetes Prices valid from April 1, 2025

Dedicated Core for Managed Kubernetes

CPU Type	Price per hour	Price group
Dedicated Core	£0.04	PG J

RAM	Price per hour	Price group
1 GB	£0.0045	PG J

Cloud Savings Plans

The following hourly rates apply to resources covered by a 1 or 3-year Savings Plan.

Resource	Term	Price per hour	Price group
1 Dedicated Core	1 Year	£0.0290	PG C1
1 Dedicated Core	3 Years	£0.0204	PG D
1 GB RAM	1 Year	£0.0032	PG C1
1 GB RAM	3 Years	£0.0023	PG D

Please note: Subscription to the Cloud Savings Plans is done through the DCD, subject to the acceptance of the Purchase Conditions.

vCPU for Managed Kubernetes

CPU Type	Price per hour	Price group
vCPU	£0.016	PG C1

RAM	Price per hour	Price group
1 GB	£0.0023	PG C1

Block Storage for Managed Kubernetes

	Price per 30 days/GB	Price group
HDD	£0.04	PG J
SSD Premium	£0.14	PG C2

Private Container Registry

Storage	Price per 30 days	Price group
1 GB	£0.04	PG C1
Vulnerability Scanning	Price per 30 days	Price group
-	-	-
per 1 GB	£0.018	PG C1

Database as a Service

In-Memory DB

Core	Price per hour	Price group
1	£0.054	PG H

RAM	Price per hour	Price group
1 GB	£0.006	PG H

Storage	Price per 30 days	Price group
1 GB SSD Premium	£0.14	PG C2

MariaDB

Core	Price per hour	Price group
1	£0.054	PG H

RAM	Price per hour	Price group
1 GB	£0.006	PG H

Storage	Price per 30 days	Price group
1 GB SSD Premium	£0.14	PG C2
1 GB Backup	£0.0063	PG C1

MongoDB

Playground Edition

Instance Type	vCPU / RAM / Storage	Price per hour	Price group
first instance	1 / 2 GB /50 GB	no charge	-
further instances	1/ 2 GB / 50 GB	£0.035	PG C1

Business Edition

Instance Type	vCPU / RAM / Storage	Price per hour	Price group
XS	1 / 2 GB / 50 GB	£0.035	PG C1
S	2 / 4 GB / 80 GB	£0.070	PG C1
M	4 / 8 GB / 160 GB	£0.140	PG C1
L	6 / 16 GB / 320 GB	£0.280	PG C1
XL	8 / 32 GB / 640 GB	£0.560	PG C1
XXL	16 / 64 GB / 960 GB	£1.130	PG C1
4XL	32 / 128 GB / 1280 GB	£2.270	PG C1
4XL_S	32 / 128 GB / 2000 GB	£2.470	PG C1

Enterprise Edition

Core	Price per hour	Price group
1	£0.054	PG C1

RAM	Price per hour	Price group
1 GB	£0.0318	PG C1

Storage	Price per 30 days	Price group
1 GB SSD Premium	£0.14	PG C2
1 GB SSD	0.063	PG C2
1 GB HDD	£0.04	PG C1

MongoDB Backup

Backup	Price per 30 days	Price group
1 GB Backup	£0.0063	PG C1

PostgreSQL

Core	Price per hour	Price group
1	£0.06	PG H

RAM	Price per hour	Price group
1 GB	£0.007	PG H

Storage	Price per 30 days	Price group
1 GB SSD Premium	£0.14	PG C2
1 GB SSD	£0.063	PG C2
1 GB HDD	£0.04	PG H
1 GB Backup	£0.0063	PG C1

Data Analytics

Event Streams for Apache Kafka

Cluster size	Price per hour	Price group
XS	£0.29	PG H
S	£0.46	PG H
M	£0.60	PG H
L	£1.2	PG H
XL	£2.3	PG H

Network services

Data transfer

Traffic type	Volume per month	Price per 30 days/GB	Price group
Outgoing public traffic	up to 2 TB (≤ 2 TB)	no charge	PG J
Outgoing public traffic	next 8 TB (> 2 TB ≤ 10 TB)	£0.030	PG J
Outgoing public traffic	next 40 TB(> 10 TB ≤ 50 TB)	£0.025	PG J
Outgoing public traffic	next 100 TB(> 50 TB ≤ 150 TB)	£0.020	PG J
Outgoing public traffic	greater than 150 TB (> 150 TB)	£0.015	PG J
Outgoing public traffic	IONOS Public Cloud national and local	no charge	
Internal private traffic	-	no charge	
Incoming traffic	-	no charge	

IP Type	Price per 30 days	Price group
Reserved public IPv4	£5.00	PG J

Please note:

1. 1 TB equals 1,024 GB
2. Tiered pricing applies to all outgoing traffic from all IONOS Public Cloud services within the same contract, and the total cumulative amount is listed on the invoice.
3. Outgoing public traffic is classified in the following way:

3.1. IONOS Public Cloud local traffic is exchanged between services via public networks that terminate in the same IONOS Public Cloud data centre;

3.2. IONOS Public Cloud national traffic is exchanged between services via public networks that terminate in a IONOS Public Cloud data centre in the same country;

3.3. Outgoing data transfer between IONOS Public Cloud data centres in different countries, to IONOS Private Cloud, to IONOS Bare Metal Cloud as well as with the public internet, is charged as regular outgoing public traffic.

4. Internal private traffic is exchanged between services within the same IONOS Public Cloud data centre via private networks.
5. IONOS CLOUD Object Storage:

Hosted across different locations, with contract-owned buckets and user-owned buckets options.

5.1 Traffic from a VM to Object Storage is subject to the same definition of local, national, and regular public traffic for cost invoicing:

5.1.1. Outgoing VM traffic to Object Storage in the same data centre location is measured as local traffic;
5.1.2. Outgoing VM traffic to Object Storage in the same country, but to a different data centre, is measured as national traffic;
5.1.3. Outgoing VM traffic to Object Storage in data centres in different countries is charged as regular public traffic.

5.2. Outgoing Object Storage traffic from **contract-owned buckets** is charged as regular public traffic, except for data transfer to services in the same data centre or to IONOS Public Cloud data centres in the same country.

5.3. Outgoing Object Storage traffic from **user-owned buckets** is always charged as regular public traffic, and the definitions of local and national traffic do not apply, even when

this traffic is addressed to another IONOS Public Cloud service, regardless of its deployment location.

5.4. While inter-bucket data transfer is subject to charges, replication traffic both within the same region and across different regions is cost-free.

5.5. Outgoing Object Storage traffic to IONOS Bare Metal Cloud or IONOS Private Cloud is always charged as regular public traffic.

Cloud DNS

	Price per 30 days	Price group
Public Zone (Primary)	£1.35/zone	PG J
Public Zone (Secondary)	£1.35/zone	PG J
Reverse DNS record (PTR)	no charge	

Content Delivery Network

	Price per 30 days	Price group
1 CDN	£89.00	PG A
Web Application Firewall	£0.80/routing rule	PG A

Managed Application Load Balancer

	Price per hour per rule	Price group
The first 5 forwarding rules	£0.02	PG J
Each additional forwarding rule	£0.01	PG J
Processing fee for incoming and outgoing data	no charge	

Please note: A minimum charge equivalent to one rule applies, even if none are configured.

Managed NAT Gateway

	Price per hour	Price group
Basis fee	£0.04	PG J
Processing fee for incoming and outgoing data	no charge	

Managed Network Load Balancer

	Price per hour per rule	Price group
The first 5 forwarding rules	£0.02	PG J
Each additional forwarding rule	£0.01	PG J
Processing fee for incoming and outgoing data	no charge	

Please note: A minimum charge equivalent to one rule applies, even if none are configured.

VPN Gateway

VPN Gateway type	Price per hour	Price group
Standard	£0.03	PG D
Standard (High Availability)	£0.06	PG D
Enhanced	£0.07	PG D
Enhanced (High Availability)	£0.14	PG D
Premium	£0.18	PG D
Premium (High Availability)	£0.36	PG D

Please note: price per provisioned Gateway. Costs incurred for data transmission are not included, regular traffic charges apply.

Observability

Logging Service

Component	Price	Price group
per Pipeline per hour	£0.016	PG A
per GB of storage per 30 days	£0.014	PG A

Monitoring Service

	Price	Price group
Per 1 Million samples	£0.14	PG J
per GB of storage per 30 days	£0.014	PG C1

Sovereign Workplace

Nextcloud Workspace

	Price per User / 30 days	Price group
Nextcloud Workspace (1 to 24 Users)	£6.43	PG G
Nextcloud Workspace (25+ Users)	£5.43	PG G

Please note: Pay-As-You-Go pricing tiers are calculated independently for each individual Workspace. Each Workspace supports up to 200 users.

Cloud Savings Plans

The following rate applies to user seats covered by a 1-year Savings Plan. Savings Plans require a minimum of 25 users per plan and are applied at the contract level, meaning your plan's seats can be flexibly distributed across multiple Workspaces.

	Term	Price per User / 30 days	Price group
Nextcloud Workspace	1 Year	£4.75	PG E

Private Cloud powered by VMware

Basic bundle (Cluster consisting of 3 hosts, incl. VMware licences)	Price per month	Price group
Silver-48	£3,395.00	PG 1p
Silver-96	£4,600.00	PG 1p
Gold-192 v2	£5,380.00	PG 1p
Gold-512	£7,385.00	PG 1p
M-512	£5,400.00	PG 1p
L-768	£8,000.00	PG 1p
XL-1024	£10,900.00	PG 1p

Additional host (incl. VMware licences)	Price per month	Price group
Silver-48	£740.00	PG 1p
Silver-96	£915.00	PG 1p
Gold-192 v2	£1,350.00	PG 1p
Gold-512	£2,220.00	PG 1p
M-512	£1,400.00	PG 1p
L-768	£2,300.00	PG 1p
XL-1024	£3,100.00	PG 1p

Traffic type	Price	Price group
Outgoing	no charge	PG 1p
Incoming	no charge	PG 1p
National	no charge	PG 1p

Backup	Price	Price group
Cloud Backup	£/GB 0.07	PG 1p

Additional IP address(es)	Price per month	Price group
1 Additional IP address	£5.00	PG 1p
Subnet/28 (16)	£24.00	PG 1p
Subnet/27 (32)	£32.00	PG 1p

Network services	Price per month	Price group
IPSEC VPN	£50.00	PG 1p

Additional storage	Price per month	Price group
NFS Datastore (min. 2 TB)	£72.00/TB	PG 1p

VCDA Protection	Price per month	Price group
VCDA Protection	£50.00/VM	PG 1p

Please note: Only VM’s that are protected will be invoiced, the migration feature can be used totally free.

Change hardware model	Price (One-Time Fee)
Change hardware model	£300.00

Please note: Outgoing traffic to IONOS Bare Metal Cloud and Private Cloud powered by VMware is recorded as regular public traffic.

v2026-05-13, prices valid from 13/05/2026

IONOS CLOUD Limited, Discovery House, 154 Southgate Street, Gloucester, GL1 2EX, United Kingdom Registered in England and Wales, Company number 3953678 - VAT No GB 752539027

IONOS SE (German)

Das Angebot der IONOS SE für die Produkte der IONOS CLOUD richtet sich ausschließlich an Gewerbetreibende, nicht an Verbraucher. Preise sind Nettopreise und verstehen sich zuzüglich etwaig anfallender gesetzlicher Umsatzsteuer.

Als IONOS CLOUD-Kunde profitieren Sie von einem technischen 24/7-Support per Telefon und E-Mail.

AI

AI Model Hub

Large Language Models	Preis pro Million Input-Token	Preis pro Million Output-Token	Preisgruppe
Llama 3.1 8B Instruct	0,15 €	0,15 €	PG A
Mistral Nemo Instruct	0,15 €	0,15 €	PG A
Mistral Small 24B Instruct	0,1 €	0,3 €	PG A
gpt-oss-120b	0,15 €	0,65 €	PG A
Llama 3.3 70B Instruct	0,65 €	0,65 €	PG A
Llama 3.1 405B Instruct	1,75 €	1,75 €	PG A

Code-Modelle	Preis pro Million Input-Token	Preis pro Million Output-Token	Preisgruppe
Qwen3-Coder-Next (80B)	0,15 €	0,80 €	PG A
Code Llama 13B instruct HF	0,45 €	0,45 €	PG A

Bitte beachten Sie: Code Llama 13B instruct HF wird voraussichtlich am 21. Mai 2026 eingestellt.

Vision-Language-Modelle	Preis pro Million Input-Token	Preis pro Million Output-Token	Preisgruppe
LightOnOCR 2	0,15 €	0,30 €	PG A

Text-to-Image-Modelle	Preis pro Bild	Preisgruppe
FLUX.1 [schnell]	0,0288 €	PG A

Embedding-Modelle	Preis pro Million Token	Preisgruppe
paraphrase-multilingual-mpnet-base-v2	0,01 €	PG A
bge-large-en-v1.5	0,015 €	PG A
bge-m3	0,02 €	PG A
Qwen3-VL-Embedding-8B	0,10 €	PG A
Qwen3-VL-Reranker-8B	0,04 €	PG A

Document Collections (Legacy)

Vektordatenbank	Preis pro Million Token	Preisgruppe
Speicher	0,01 €/30 Tage	PG A

Bitte beachten Sie: Die Funktion „Document Collections“ steht für neue Deployments nicht mehr zur Verfügung.

Backup and Storage

Backup Service

Backup Storage Modi	Preis pro 30 Tage	Preisgruppe
Cloud Storage für Orphaned Backups	0,07 €/GB	PG F
Cloud Storage für Windows oder Mac WS Backup	0,07 €/GB	PG F
Cloud Storage für Windows oder Linux Server Backup	0,07 €/GB	PG F
Cloud Storage für Virtual Machine Backup	0,07 €/GB	PG F
Backup Storage auf lokalem Datenträger	0,07 €/GB	PG F
Servicegebühr für Backup Storage im IONOS CLOUD Object Storage	0,016 €/GB	PG F

Bitte beachten Sie: Für die Speicherung von Backups im IONOS CLOUD Object Storage ist das Direct Backup to Public Cloud Feature Pack erforderlich.

Acronis Advanced Features	Preis pro 30 Tage	Preisgruppe
Advanced Security Feature Paket (pro registrierter VM, Server, Workstation)	1,68 €	PG A
Advanced Management Feature Paket (pro registrierter VM, server, Workstation)	1,68 €	PG A
Direct Backup to Public Cloud Feature Paket (pro registrierter VM, Server, Workstation)	5,88 €	PG A

Block Storage

Block Storage	Preis pro 30 Tage	Preisgruppe
HDD	0,04 €/GB	PG J
SSD Standard	0,07€/GB	PG C2
SSD Premium	0,15€/GB	PG C2
Snapshot	0,04 €/GB	PG J

Kostenlose Storage IOPS (Input/Output Operationen pro Sekunde)

IONOS Cloud Object Storage

Storage	Preis pro 30 Tage	Preisgruppe
1GB	0,007 €	PG C1

API Call Typen	Preis pro 10.000 Calls	Preisgruppe
PUT/COPY/POST/LIST	kostenfrei	
GET	kostenfrei	
DELETE	kostenfrei	

Data Storage Zugriff	Preis pro GB	Preisgruppe
Write	kostenfrei	
Read	kostenfrei	

- 1 Gigabyte (GB) entspricht 1024 Megabytes (MB).
- Der Speicherplatz wird mit einer Granularität von einer Stunde abgerechnet.

Bitte beachten Sie: Die Kosten für die Datenübertragung sind nicht im Angebot enthalten und werden gemäß den aktuell gültigen Preisen addiert.

Network File Storage

Storage	Preis pro 30 Tage	Preisgruppe
1 TB	154 €	PG B

Bitte beachten Sie: Die Mindestgröße für die Bereitstellung beträgt 2 TB, die Höchstgröße 42 TB.

Compute

Compute Engine

vCPU-Server

CPU Typ	Preis pro Stunde	Preisgruppe
vCPU	0,012€	PG C1

RAM	Preis pro Stunde	Preisgruppe
1 GB	0,002€	PG C1

Dedicated Core-Server

Dedicated Core - CPU Typ	Preis pro Stunde	Preisgruppe
AMD Opteron*	0,02 €	PG J
AMD EPYC Milan	0,036 €	PG J
AMD EPYC Turin	0,042 €	PG J
Intel® Core™ Haswell	0,04 €	PG J
Intel® Core™ Skylake	0,04 €	PG J
Intel® Core™ Ice Lake	0,04 €	PG J
Intel® Core™ Sierra Forest	0,046 €	PG J

RAM	Preis pro Stunde	Preisgruppe
1 GB	0,0045 €	PG J

**Der AMD Opteron ist für neue Bereitstellungen nicht verfügbar.*

Cloud Savings Plans

Die folgenden stündlichen Raten gelten für Ressourcen, die durch einen 1-jährigen oder 3-jährigen Savings Plan abgedeckt sind.

Ressource	Laufzeit	Preis pro Stunde	Preisgruppe
1 Dedicated Core	1 Jahr	0,0340 €	PG C1
1 Dedicated Core	3 Jahre	0,0240 €	PG D
1 GB RAM	1 Jahr	0,0038 €	PG C1
1 GB RAM	3 Jahre	0,0027 €	PG D

Bitte beachten Sie:

1. Die Savings Plans decken alle Dedicated Core CPU-Typen ab, außer AMD Opteron.
2. Die Buchung der Cloud Savings Plans erfolgt über den DCD und unterliegt der Annahme der Einkaufsbedingungen.

Cloud GPU VM

Template	GPU Model	GPU Type	Number of GPUs	Dedicated CPUs	RAM (GiB)	NVMe Storage (GB)	Price per hour	Price group
H200-S	NVIDIA H200	H200 PCIe	1	15	267	1024	3,00€	PG C1
H200-M	NVIDIA H200	H200 PCIe	2	30	534	1536	6,00€	PG C1
H200-L	NVIDIA H200	H200 PCIe	4	60	1068	2048	12,00€	PG C1
H200-XL	NVIDIA H200	H200 PCIe	8	127	2136	4096	24,00€	PG C1

Microsoft Lizenzen

Microsoft Windows Server (2016, 2019, 2022, 2025)

CPU Typ	Preis pro Stunde	Preisgruppe
Pro vCPU	0,0149 €	PG A
Pro AMD Core	0,0299 €	PG A
Pro Intel® Core™	0,0299 €	PG A

Microsoft SQL Server

Edition	Preis pro Stunde	Preisgruppe
Web Edition	0,0181 €	PG A
Server Standard	0,25 €	PG A
Enterprise Edition	0,99 €	PG A

Die Microsoft SQL Server Lizenzpakete beinhalten die Nutzung von 2 Cores bzw. 1 Core-Paar. Pro Server werden immer mindestens 2 Lizenzpakete abgerechnet.

RedHat Enterprise Linux

RHEL on vCPU/Dedicated Core	Preis pro Stunde	Preisgruppe
RHEL Server Small Virtual Node (1-8 vCPUs/ Cores)	0,0143 € pro vCPU oder Core	PG A
RHEL Server Medium Virtual Node (9-127 vCPUs/ Cores)	0,0107 € pro vCPU oder Core	PG A
RHEL Server Large Virtual Node (128 or more vCPUs/ Cores)	0,0095 € pro vCPU oder Core	PG A

IONOS CLOUD Cubes

Verfügbar ab 1. Oktober 2024

Konfiguration	vCPU / RAM / Storage	Preis pro Stunde	Preisgruppe
Basic Cube XS	1 / 2 GB / 60 GB	0,007 €	PG C1
Basic Cube S	2 / 4 GB / 120 GB	0,013 €	PG C1
Basic Cube M	4 / 8 GB / 240 GB	0,024 €	PG C1
Basic Cube L	8 / 16 GB / 480 GB	0,044 €	PG C1
Basic Cube XL	16 / 32 GB / 960 GB	0,083 €	PG C1
Memory Cube S	2 / 8 GB / 120 GB	0,017 €	PG C1
Memory Cube M	4 / 16 GB/ 240 GB	0,032 €	PG C1
Memory Cube L	8 / 32 GB / 480 GB	0,061 €	PG C1
Memory Cube XL	16 / 64 GB / 960 GB	0,114 €	PG C1

Verfügbar bis 30. September 2024

Konfiguration	vCPU / RAM / Storage	Preis pro Stunde	Preisgruppe
Cubes XS	1 / 1 GB / 30 GB	0,0028 €	PG C1
Cubes S	1 / 2 GB / 50 GB	0,0056 €	PG C1
Cubes M	2 / 4 GB / 80 GB	0,01 €	PG C1
Cubes L	4 / 8 GB / 160 GB	0,021 €	PG C1
Cubes XL	6 / 16 GB / 320 GB	0,039 €	PG C1
Cubes XXL	8 / 32 GB / 640 GB	0,083 €	PG C1
Cubes 3XL	12 / 48 GB/ 960 GB	0,125 €	PG C1
Cubes 4XL	16 / 64 GB / 1280 GB	0,167 €	PG C1

Microsoft Lizenzen

Microsoft Windows Server (2016, 2019, 2022, 2025)

Lizenz für Cubes	Preis pro Stunde	Preisgruppe
pro vCPU	0,0069 €	PG A

RedHat Enterprise Linux

RHEL on Cubes	Preis pro Stunde	Preisgruppe
RHEL Server Small Virtual Node (1-8 vCPUs/ Cores)	0,0143 € pro vCPU oder Core	PG A
RHEL Server Medium Virtual Node (9-127 vCPUs/ Cores)	0,0107 € pro vCPU oder Core	PG A
RHEL Server Large Virtual Node (128 or more vCPUs/ Cores)	0,0095 € pro vCPU oder Core	PG A

Containers

Managed Kubernetes

Neue Managed Kubernetes Preise gültig ab dem 1. April 2025

Dedicated Core für Managed Kubernetes

CPU-Typ	Preis pro Stunde	Preisgruppe
Dedicated Core	0,04€	PG J

RAM	Preis pro Stunde	Preisgruppe
1 GB	0,0045€	PG J

Cloud Savings Plans

Die folgenden stündlichen Raten gelten für Ressourcen, die durch einen 1-jährigen oder 3-jährigen Savings Plan abgedeckt sind.

Ressource	Laufzeit	Preis pro Stunde	Preisgruppe
1 Dedicated Core	1 Jahr	0,0340 €	PG C1
1 Dedicated Core	3 Jahre	0,0240 €	PG D
1 GB RAM	1 Jahr	0,0038 €	PG C1
1 GB RAM	3 Jahre	0,0027 €	PG D

Bitte beachten Sie: Die Buchung der Cloud Savings Plans erfolgt über den DCD und unterliegt der Annahme der Einkaufsbedingungen.

vCPU für Managed Kubernetes

CPU-Typ	Preis pro Stunde	Preisgruppe
vCPU	0,017€	PG C1

RAM	Preis pro Stunde	Preisgruppe
1 GB	0,0025€	PG C1

Block-Storage für Managed Kubernetes

	Preis pro 30 Tage/GB	Preisgruppe
HDD	0,04€	PG J
SSD Premium	0,15€	PG C2

Private Container Registry

Storage	Preis pro 30 Tage	Preisgruppe
1 GB	0,04 €	PG C1

Vulnerability Scanning	Preis pro 30 Tage	Preisgruppe
Pro 1 GB	0,02 €	PG C1

Database as a Service

In-Memory DB

Core	Preis pro Stunde	Preisgruppe
1	0,06 €	PG H

RAM	Preis pro Stunde	Preisgruppe
1 GB	0,007 €	PG H

Storage	Preis pro 30 Tage	Preisgruppe
1 GB SSD Premium	0,15 €	PG C2

MariaDB

Core	Preis pro Stunde	Preisgruppe
1	0,06 €	PG H

RAM	Preis pro Stunde	Preisgruppe
1 GB	0,007 €	PG H

Storage	Preis pro 30 Tage	Preisgruppe
1 GB SSD Premium	0,15 €	PG C2
1 GB Backup	0,007 €	PG C1

MongoDB

Playground Edition

Instanz Typ	vCPU / RAM / Storage	Preis pro Stunde	Preisgruppe
Erste Instanz	1 / 2 GB /50 GB	kostenfrei	-
Weitere Instanzen	1/ 2 GB / 50 GB	0,035 €	PG C1

Business Edition

Instanz Typ	vCPU / RAM / Storage	Preis pro Stunde	Preisgruppe
XS	1 / 2 GB / 50 GB	0,035 €	PG C1
S	2 / 4 GB / 80 GB	0,070 €	PG C1
M	4 / 8 GB / 160 GB	0,140 €	PG C1
L	6 / 16 GB / 320 GB	0,280 €	PG C1
XL	8 / 32 GB / 640 GB	0,560 €	PG C1
XXL	16 / 64 GB / 960 GB	1,130 €	PG C1
4XL	32 / 128 GB / 1280 GB	2,270 €	PG C1
4XL_S	32 / 128 GB / 2000 GB	2,470 €	PG C1

Enterprise Edition

Core	Preis pro Stunde	Preisgruppe
1	0,06 €	PG C1

RAM	Preis pro Stunde	Preisgruppe
1 GB	0,0353 €	PG C1

Storage	Preis pro 30 Tage	Preisgruppe
1 GB SSD Premium	0,15 €	PG C2
1 GB SSD	0,07 €	PG C2
1 GB HDD	0,04 €	PG C1

MongoDB Backup

Backup	Preis pro 30 Tage	Preisgruppe
1 GB Backup	0,007 €	PG C1

PostgreSQL

Core	Preis pro Stunde	Preisgruppe
1	0,06 €	PG H

RAM	Preis pro Stunde	Preisgruppe
1 GB	0,007 €	PG H

Storage	Preis pro 30 Tage	Preisgruppe
1 GB SSD Premium	0,15 €	PG C2
1 GB SSD	0,07 €	PG C2
1 GB HDD	0,04 €	PG H
1 GB Backup	0,007 €	PG C1

Data Analytics

Event Streams for Apache Kafka

Cluster size	Preis pro Stunde	Preisgruppe
XS	0,32 €	PG H
S	0,51 €	PG H
M	0,67 €	PG H
L	1,3 €	PG H
XL	2,5 €	PG H

Network services

Datenübertragung

Typ des Traffics	Volumen pro Monat	Preis	Preisgruppe
Ausgehender öffentlicher Traffic	für die ersten 2 TB (≤ 2 TB)	kostenfrei	PG J
Ausgehender öffentlicher Traffic	die nächsten 8 TB (> 2 TB ≤ 10 TB)	0,030 €/GB	PG J
Ausgehender öffentlicher Traffic	die nächsten 40 TB (> 10 TB ≤ 50 TB)	0,025 €/GB	PG J
Ausgehender öffentlicher Traffic	die nächsten 100 TB (> 50 TB ≤ 150 TB)	0,020 €/GB	PG J
Ausgehender öffentlicher Traffic	über 150 TB (> 150 TB)	0,015 €/GB	PG J
Ausgehender öffentlicher Traffic	IONOS Public Cloud national und lokal	kostenfrei	
Interner privater Traffic	-	kostenfrei	
Eingehender Traffic	-	kostenfrei	

IP Typ	Preis pro 30 Tage	Preisgruppe
Reservierung öffentliche IPv4 Adresse	5,00 €	PG J

Bitte beachten Sie:

1. 1 TB entspricht 1,024 GB
2. Die Preisstaffelung gilt für den gesamten ausgehenden Datenverkehr aller IONOS Public Cloud Services innerhalb desselben Vertrags. Der kumulierte Gesamtbetrag wird auf der Rechnung aufgeführt.
3. Ausgehender öffentlicher Datenverkehr wird wie folgt klassifiziert:

3.1. Der lokale IONOS Public Cloud-Traffic wird zwischen den Diensten über öffentliche Netze ausgetauscht, die im selben IONOS Public Cloud-Rechenzentrum enden;

3.2. Der nationale IONOS Public Cloud-Traffic wird zwischen Diensten über öffentliche Netze ausgetauscht, die in einem IONOS Public Cloud-Rechenzentrum im selben Land enden;

3.3. Ausgehender Datenverkehr zwischen IONOS Public Cloud-Rechenzentren in verschiedenen Ländern, zu IONOS Private Cloud, zu IONOS Bare Metal Cloud sowie mit dem öffentlichen Internet wird als regulärer ausgehender öffentlicher Trafik berechnet.

5. Interner privater Datenverkehr wird zwischen Services innerhalb desselben IONOS Public Cloud Rechenzentrums über private Netzwerke ausgetauscht.
6. IONOS CLOUD Object Storage:

Gehostet an verschiedenen Standorten, mit Optionen für vertragseigene Buckets und benutzereigene Buckets.

5.1 Der Traffic von einer VM zu Object Storage unterliegt der gleichen Definition von lokalem, nationalem und regulärem öffentlichen Traffic für die Kostenberechnung:

5.1.1. Ausgehender VM-Datenverkehr zu Object Storage am gleichen Rechenzentrumsstandort wird als lokaler Trafik gewertet; 5.1.2. Ausgehender VM-Datenverkehr zu Object Storage im selben Land, aber zu einem anderen Rechenzentrum, wird als nationaler Verkehr gemessen;
5.1.3. Ausgehender VM-Datenverkehr zu Object Storage in Rechenzentren in verschiedenen Ländern wird als regulärer öffentlicher Verkehr berechnet.

5.2. Ausgehender Object Storage-Traffic von **vertragseigenen Buckets** wird als regulärer öffentlicher Datenverkehr berechnet, mit Ausnahme von Datentransfers zu Diensten im selben Rechenzentrum oder zu IONOS Public Cloud-Rechenzentren im selben Land.

5.3. Ausgehender Object Storage-Traffic von **benutzereigenen Buckets** wird immer als regulärer öffentlicher Verkehr berechnet, und die Definitionen von lokalem und nationalem Verkehr gelten nicht, auch wenn dieser Verkehr an einen anderen IONOS Public Cloud-Service gerichtet ist, unabhängig von dessen Bereitstellungsort.

5.4 Während für den Datentransfer zwischen Buckets Gebühren anfallen, ist der Replikationsdatenverkehr sowohl innerhalb der gleichen Region als auch zwischen verschiedenen Regionen kostenlos.

5.5. Ausgehender Object Storage-Verkehr zu IONOS Bare Metal Cloud oder IONOS Private Cloud wird immer als regulärer öffentlicher Datenverkehr berechnet.

Cloud DNS

	Preis pro 30 Tage	Preisgruppe
Public Zone (Primary)	1,50 €/Zone	PG J
Public Zone (Secondary)	1,50 €/Zone	PG J
Reverse DNS-Eintrag (PTR)	kostenfrei	

Content Delivery Network

	Preis pro 30 Tage	Preisgruppe
1 CDN	99,00 €	PG A
Web Application Firewall	0,90 €/routing rule	PG A

Managed Application Load Balancer

	Preis pro Stunde pro Regel	Preisgruppe
Die ersten 5 Weiterleitungsregeln	0,02 €	PG J
Pro zusätzlicher Weiterleitungsregel	0,01 €	PG J
Verarbeitungsgebühr für ein- und ausgehende Daten	kostenfrei	

Bitte beachten Sie: Es gilt eine Mindestgebühr in Höhe einer Regel, auch wenn keine konfiguriert ist.

Managed NAT Gateway

	Preis pro Stunde	Preisgruppe
Basisgebühr	0,04 €	PG J
Verarbeitungsgebühr für ein- und ausgehende Daten	kostenfrei	

Managed Network Load Balancer

	Preis pro Stunde pro Regel	Preisgruppe
Die ersten 5 Weiterleitungsregeln	0,02 €	PG J
Pro zusätzlicher Weiterleitungsregel	0,01 €	PG J
Verarbeitungsgebühr für ein- und ausgehende Daten	kostenfrei	

Bitte beachten Sie: Es gilt eine Mindestgebühr in Höhe einer Regel, auch wenn keine konfiguriert ist.

VPN Gateway

VPN Gateway-Typ	Preis pro Stunde	Preisgruppe
Standard	0,032 €	PG D
Standard (High Availability)	0,064 €	PG D
Enhanced	0,076 €	PG D
Enhanced (High Availability)	0,15 €	PG D
Premium	0,20 €	PG D
Premium (High Availability)	0,40 €	PG D

Bitte beachten Sie: Preis pro bereitgestelltem Gateway. Die Kosten für die Datenübertragung sind nicht inbegriffen, es fallen die üblichen Verkehrsgebühren an.

Observability

Logging Service

Komponente	Preis	Preisgruppe
Pro "Pipeline" pro Stunde	0,018 €	PG A
Pro GB Speicher pro 30 Tage	0,015 €	PG A

Monitoring Service

	Preis	Preisgruppe
Pro 1 Million Stichproben	0,15 €	PG J
Pro GB Speicher pro 30 Tage	0,015 €	PG C1

Sovereign Workplace

Nextcloud Workspace

	Preis pro Anwender / 30 Tage	Preisgruppe
Nextcloud Workspace (1 bis 24 Anwender)	7,56 €	PG G
Nextcloud Workspace (25 bis 200 Anwendern)	6,39 €	PG G

Bitte beachten Sie: Die Preisstaffeln im Pay-As-You-Go-Modell werden für jeden Workspace unabhängig berechnet. Pro Workspace sind maximal 200 Anwender möglich.

Cloud Savings Plans

Der folgende Tarif gilt für Anwenderplätze, die durch einen 1-Jahres-Savings-Plan abgedeckt sind. Savings Plans erfordern ein Minimum von 25 Anwendern pro Plan und gelten auf Vertragsebene. Das bedeutet, dass die Plätze Ihres Plans flexibel auf mehrere Workspaces verteilt werden können.

	Laufzeit	Preis pro Anwender / 30 Tage	Preisgruppe
Nextcloud Workspace	1 Jahr	5,59 €	PG E

Private Cloud powered by VMware

Basic bundle (Cluster bestehend aus 3 Hosts, incl. VMware licenses)	Preis pro Monat	Preisgruppe
Silver-48	3.900,00 €	PG 1p
Silver-96	5.280,00 €	PG 1p
Gold-192 v2	6.180,00 €	PG 1p
Gold-512	8.480,00 €	PG 1p
M-512	6.300,00 €	PG 1p
L-768	9.300,00 €	PG 1p
XL-1024	12.700,00 €	PG 1p

Zusätzlicher Host	Preis pro Monat	Preisgruppe
Silver-48	850,00 €	PG 1p
Silver-96	1.050,00 €	PG 1p
Gold-192 v2	1.550,00 €	PG 1p
Gold-512	2.550,00 €	PG 1p
M-512	1.600,00 €	PG 1p
L-768	2.600,00 €	PG 1p
XL-1024	3.700,00 €	PG 1p

Traffic Typ	Preis	Preisgruppe
Outgoing	kostenfrei	PG 1p
Incoming	kostenfrei	PG 1p
National	kostenfrei	PG 1p

Backup	Preis pro Monat	Preisgruppe
Cloud Backup	0,07 €/GB	PG 1p

Zusätzliche IP-Adresse(n)	Preis pro Monat	Preisgruppe
1 zusätzliche IP Adresse	5,00 €	PG 1p
Subnet/28 (16)	24,00 €	PG 1p
Subnet/27 (32)	32,00 €	PG 1p

Network services	Preis pro Monat	Preisgruppe
IPSEC VPN	50,00 €	PG 1p

Zusätzlicher Storage	Preis pro Monat	Preisgruppe
NFS Datastore (min. 2 TB)	72,00 €/TB	PG 1p

VCDA Protection	Preis pro Monat	Preisgruppe
VCDA Protection	50,00 €/VM	PG 1p
Bitte beachten:		
Nur geschützte VMs werden berechnet, die Migrationsfunktion kann vollkommen kostenlos genutzt werden.		

Hardwaremodell ändern	Preis (einmalig)
Hardwaremodell ändern	300,00 €

Bitte beachten: Ausgehender Traffic zur IONOS Bare Metal Cloud und Private Cloud powered by VMware wird als regulärer öffentlicher Traffic erfasst.

v2026-05-13, Preise gültig ab 13.05.2026

4. FAQ

4.1. Overview

IONOS CLOUD is a cloud infrastructure platform offering compute, storage, networking, and managed services across European and US data centres. This FAQ covers the most common questions from customers about getting started, managing accounts, understanding billing, and using key services.

If your question is not listed here, contact the IONOS CLOUD support team directly—see **Contact IONOS CLOUD Support.**

Browse by category

- Account Settings ›
- Compute Engine ›
- Costs and Billing ›
- Data Center Designer ›
- General Inquiries ›

Payment Options ›

Trial Conditions ›

Most asked questions

How do I recover my IONOS CLOUD password?

Go to the **IONOS CLOUD login page**, click **Forgot password**, and enter your registered email address. You will receive a reset link within a few minutes. For more detail, see **Account Settings** —Password Recovery.

How is IONOS CLOUD billing calculated?

IONOS CLOUD uses per-minute billing for most resources. Costs are calculated based on actual usage and consolidated into a monthly invoice. You can view real-time cost estimates in the DCD and set spending limits to control expenses. See **Costs and Billing** for the full breakdown.

What data centre locations does IONOS CLOUD offer?

IONOS CLOUD operates data centres across Europe and the United States, including Frankfurt, Berlin, London, Paris, Newark, Las Vegas, Logroño, and others. Each resource is assigned to a location at creation and cannot be moved after provisioning. See — **Data Center Designer** locations.

Does IONOS CLOUD offer a free trial?

Yes. IONOS CLOUD offers a trial account with credits to test the platform at no cost. The trial period allows access to core services, including virtual servers, storage, and networking. See **Trial conditions** for eligibility and resource limits.

Where can I download my invoices?

Invoices are available in the IONOS CLOUD Control Panel under **Billing** > **Invoices**. They are generated monthly and available as PDF downloads. See **Costs and Billing**—Invoices.

How do I manage user permissions in the DCD?

User permissions in the Data Center Designer are managed at the contract level. An administrator can assign roles and resource-level access from the **Users and Groups** section. See **Data Center Designer** —User permissions.

Is IONOS CLOUD compliant with GDPR and ISO standards?

IONOS CLOUD maintains compliance with GDPR, ISO 27001, ISO 9001, and other certifications. All data stored in European data centers remains within the EU. See **General Inquiries** —Compliance for full certification details.

Related resources

- **Service Catalog** —All available IONOS CLOUD services.
- **Prices** - Detailed pricing for all resources.
- **Glossary of Terms** —Definitions of technical terms used across the documentation.
- **IONOS CLOUD Support** —SLAs and support terms.

> New to IONOS CLOUD? **Sign up** for a free trial and explore the platform with trial credits. Startup companies can also apply for the **IONOS CLOUD Startup Program** for dedicated support and offers.

Account Settings

How can I recover my password?

You can use the **Password Recovery** option to recover your password. Alternatively, you can select the **forgot password?** option on the **DCD Log** in console.

How can I update my email address?

Contract owners can contact their account manager or the support team to update the email address registered for their IONOS CLOUD accounts. We will verify the request and update the email address accordingly. All other users can update their email address for logging in by contacting the contract owner or administrator.

Do I need a verification code to log in to IONOS CLOUD account?

No verification code is required to log in to your IONOS CLOUD account for the first time. Once the **2-Step Verification** is **activated**, you cannot log in again without the code. If the **Authenticator** application is deleted or the device it's on is not available, you must contact the support team and request to turn off the **2-Step Verification**. You can then log in again and reactivate the **2-Step Verification** yourself.

How do I delete my account?

To delete your account, contact your account manager or call us at +49 30 57700 840 or write us an email at **beratung@cloud.ionos.de**. To immediately stop your billing, delete the allocated resources and data, and let us know when you want the contract to end. The data cannot be recovered if your resources have been deleted.

Compute Engine

What are cloud servers?

The cloud server is suitable for anyone who wants a robust IT infrastructure at a reasonable price. It is the perfect foundation for various web projects, whether it is a private homepage or company website, a robust framework for web developers, or a high-performance gaming server. With cloud servers, you can flexibly balance the necessary resources and pay only for what is actually used. If you are working on larger projects, you can use servers for various purposes. For example, consider you have a frontend server paired with a load balancer, with database servers managed through third-party software using an API. In such cases, cloud server solutions offer a seamless match to address these diverse scenarios.

Are cloud servers suitable for enterprise hosting?

Yes, the cloud plans provided by IONOS CLOUD are ideal for small and medium businesses. Choose the appropriate tariff for your business needs based on the required server resources. They are usually differentiated by Central Processing Unit (CPU), Random Access Memory (RAM), and memory capacity. Fast Solid State Disk (SSD) with low access times is available in every data transfer and storage plan. Whenever you need more resources, you can adjust these resources in real-time as needed. We recommend using **Compute Engine** for larger companies needing more performance and server resources.

How can I change the data center location?

You can specify the location of a data center only when you provision it for the first time. Once the data center is provisioned, you can no longer change its location.

Why can I not access my Virtual Machine (VM) through the Remote Console?

DCD Remote Console relies on Virtual Network Computing (VNC) and requires **port 5900**. Always use the latest **Java version**. Make the following ports available in your firewall for outgoing connection requests:

- **For HTTP:** port 80
- **For HTTPS:** port 443
- **For VNC:** port 5900

The Remote Console becomes available immediately once the provisioning of your server is complete.

How can I find the location of my data center?

In the **DCD**, go to the **Dashboard** > **My Data Centers**. The **Region** column has the list of the location of your data center.

Alternatively, you can use the **Start Center** option to view the location of your Virtual Data Center.

Costs and Billing

From where can I download my invoices?

You will receive monthly invoices by email to the specified invoice recipient. If you have any questions about your accounts, contact by email at **accounting@billing.ionos.com** or your IONOS CLOUD account manager.

Can I limit my costs?

Your account manager can provide you with information on cost limits.

Can I calculate my expenses in advance?

Your account manager can provide you with specific pricing information. You can also contact us at +49 30 57700 840 or by email at **beratung@cloud.ionos.de**.

How is the traffic measured and billed?

You pay only for the **outbound traffic**, which passes through the IONOS CLOUD physical gateways. External inbound and internal traffic are free of charge. We meter on **Layer-2 (Ethernet)** frames as they traverse the physical gateways of IONOS CLOUD data centers. Usage-based charges are rounded to the nearest whole cent by IONOS CLOUD metering and billing system. For more information, see **IONOS CLOUD website**.

What are the charges for traffic over public IPs?

Traffic over public IPs is considered internal traffic and is free of charge.

What are the costs for a firewall or a load balancer?

Firewalls and load balancers are provided free of charge.

How much would a data transfer between data centers cost?

Data transfers between data centers in the same geographical regions are free. Traffic exchanged between data centers in different geographical regions is charged. This means transfers between the data center in Karlsruhe and one in Frankfurt are free, while traffic between the same data center in Karlsruhe and one in Las Vegas is billed.

Do you offer discounts on contracts?

Contact your account manager to get information about contract discounts.

Does IONOS CLOUD have special offers for students, freelancers, or startups?

We have special offers for students, freelancers, and startups. For more information, see **IONOS CLOUD Startup Program**.

4.5. Data Center Designer

How do I connect to a server using the DCD?

Each server has a Remote Console that can be accessed using the DCD. This Remote Console is available in two versions, which is used depending on your browser:

- Java Applet
- HTML5

We recommend using the **HTML5** version.

How do I manage user permissions?

In the **DCD**, go to **Menu** > **Management** > **Users and Groups**. The **User Manager** will open up. It allows contract owners and administrators to organize users. From the **User Manager**, the administrator can provide users with permissions and access to the resources of an IONOS CLOUD account.

What is the difference between a shutdown and a power stop?

When a Virtual Machine (VM) is shut down only at the operating system level, the hardware resources it consumes remain allocated to the VM, and the costs for these resources continue to be charged. To avoid unnecessary costs, use the DCD to stop the VM and deallocate its resources. Costs for CPU cores and RAM will then no longer apply. Associated storage, however, is not deallocated automatically when the VM is stopped and will continue to be charged until deleted manually.

How can I fix the error message "Missing required Permissions manifest attribute in main jar"?

Add an exception for the **https://dcd.ionos.com** URL to resolve this error in your browser's Java configuration.

Why are there no switch elements?

DCD does not use switches for creating network segments because functions such as switching, routing, and forwarding are deeply integrated into the IONOS CLOUD Software Defined Network

(SDN) stack. The traffic distribution is handled by IONOS CLOUD.

How is provisioning conducted?

Provisioning allocates billable resources required to start operating your VDC. Provisioning means that the infrastructure of your virtual data center is made available for use. Once all the changes to your Virtual Data Center (VDC) have been completed, you can begin the provisioning process. Please note that this process cannot be undone.

An existing data center can be modified at any time. Note that the provisioning process cannot be undone. Your password will be requested to edit some elements as an additional security measure.

Before starting provisioning, you will be alerted about any issues or errors and can review all changes. As an additional security measure for critical changes, you must enter your password before provisioning can begin.

DCD verifies changes in the following ways before provisioning:

- Validation
- Delta

Are there provisioning validation safeguards?

DCD ensures a smooth provisioning process by checking the validity of new VDC configurations. This is done before provisioning and while creating VDCs.

- **Live validation:** Identifies errors that would block the provisioning process and provides assistance with their immediate resolution. Errors and warnings that do not stop the provisioning process or messages requesting a password confirmation are not displayed.
- **Dialog box:** The new configuration is checked for errors before starting the provisioning process. For example, you receive notifications if default element names are unchanged. The built-in help function will assist you in resolving such errors.

Select the **Delta** tab in the **Provision Data Center** to review all changes before starting the provisioning process. A categorized list of changes is displayed as follows:

- Added
- Deleted
- Updated

The authorization in DCD is based on the user privileges. Users inherit permissions and access privileges granted to the group or groups of which they are members.

How are user privileges assigned?

Permissions are associated with privileges and groups. Privileges define what users are allowed to do and to which part of the DCD they have access. The following privileges can be attributed:

Contract Owners	Administrators	Users
All users who register with IONOS CLOUD are automatically the contract owners.	Administrators have the same permissions as contract owners. However, they cannot change the payment method.	A user is any member who has been added to the contract by the contract owner.
Only contract owners have access to the **User Manager**. They can view the resource and cost overview and change the payment method of the contract. They can add or remove users from your contract, create groups, and grant permissions and access privileges.	Administrators can assign administrator privileges to other users.	A user assigned to a group, at a minimum, has read access to the resources allocated to the group. Permissions are required for some activities in the DCD.
Only contract owners have access to the **User Manager**. They can view the resource and cost overview and change the payment method of the contract. They can add or remove users from your contract, create groups, and grant permissions and access privileges.	Administrators need not be managed in groups, as they automatically have access to all resources associated with the contract.	
The contract can grant users, in their contract, administrator privileges by activating the respective checkboxes in the **User Manager**.	The administrator can grant users, in their contract, administrator privileges by activating the respective checkboxes in the **User Manager**.	

How are user groups divided?

The groups facility allows a contract owner or an administrator to:

- Define what users are allowed to do in the DCD.
- Define which resources (VDC, images, snapshots, IP blocks) they have access to.

Administrators need not be managed in groups. They automatically have access to all resources associated with the contract. There is no limit to the number of groups and users that can be created.

4.6. General Inquiries

How can I ensure my data center is properly set up?

Customers must have a basic understanding of cloud hosting platforms before using IONOS CLOUD's products and services. The best way to get started is by using IONOS CLOUD DCD tutorials. These instruction sets will help you start a fundamental data center and set up initial user access and permissions.

If you require direct assistance, contact **IONOS CLOUD help**. They will check your infrastructure to ensure that everything is correctly set up. This service is completely free of charge. IONOS CLOUD support service is available 24/7.

> **Note:** Cloud Specialists cannot perform a full migration for new customers. Let us know if this is necessary so that we can find the right partner from IONOS CLOUD partner network.

If I need a certain feature but cannot find it on the website, how do I know if IONOS CLOUD offers it?

If you already have an account with IONOS CLOUD, contact your account manager. If not, you can contact the **IONOS CLOUD help**.

Do you guarantee that my data will only remain in Germany?

It depends on you, what data, and where to store it because you choose the data center's location. All information about data protection can be found in IONOS CLOUD data protection package. If you have further questions, email at **beratung@cloud.ionos.de**.

Can I get help in setting up my Virtual Data Center?

We have extensive **IONOS CLOUD Documentation** and offer many informative **webinars**. You can also customize your own infrastructure according to your needs. If you need IT service provider support, we are happy to introduce you to a partner.

Does IONOS CLOUD offer specific certifications?

Yes, IONOS CLOUD offers specific certifications. For more information, see **Certificates** and attestations. If you need additional certificates, you can either contact IONOS CLOUD customer service or contact your account manager.

Who should I contact if I want to transfer my account to another company?

To transfer your account to another company, contact your account manager, who will help you with the transfer and tell you what confirmations and information we need.

Who should I contact if I need an Auftragsverarbeitungsvertrag (AVV)?

To initiate this operation, contact your account manager. For more information on AVV, see **Order processing agreement (AVV)**.

Are IONOS CLOUD hosting services environmentally sustainable?

Yes, IONOS CLOUD services are environmentally sustainable. For more information, see **Sustainability** at IONOS CLOUD.

Payment Options

Which payment methods are available in the DCD?

You can change your payment method in the DCD at any time. You can choose between SEPA direct debit or credit card (Visa, Mastercard, and Amex).

Why must I register with a credit card?

When registering with a credit card, we can verify that you are indeed a real individual. In addition, we will use your credit card as a payment method after the end of the testing phase. If you want to keep another means of payment, such as SEPA direct debit, you can set it up at any time in the DCD.

How do I update my payment details?

You can change your payment details in your DCD account at any time. In the **DCD**, go to the **Menu** > **Profile** > **Payment Method**. Only the contract owner can administer payment data; other registered users cannot do this.

Is it possible to pay by invoice or in advance?

To learn about payment possibilities, contact your dedicated account manager.

Trial Conditions

Why should I create a test account?

With IONOS CLOUD, every customer can test the production environment. Cloud infrastructure can be thoroughly tested in real-life conditions. If you decide to continue using the services after the end of the test, the infrastructure you set up will remain, and only the account status will change. User data saved during registration will be used for billing purposes.

What happens after I create a test account?

After you enter your company and credit card details, you will have access to a test account within a few minutes. In rare cases, if, for example, a credit card or phone number verification was not completed successfully, another phone call is needed. IONOS CLOUD sales team will contact you by phone within 24 hours on working days to discuss the next steps.

Which resources are available during the trial period?

During the trial period, 8 cores, 20 GB RAM, 600 GB SSD, and HDD, and 1 static IP address, are available by the standard. Your trial period is limited to 30 days or a maximum resource consumption of 500 euros. If you need more test resources, contact your account manager.

Can I observe my trial resource usage?

You can display the resources you are currently using in the DCD. This overview includes CPU, RAM, Storage, and IP addresses. Traffic, Windows licenses, backup, or S3 object storage are not currently displayed but can be viewed in detail using the **Billing API**.

What happens if I run out of testing resources?

As soon as your test credit is spent, your account will be converted to a productive account. Saved payment details will be used for billing purposes.

What happens after the end of the test phase?

After the test phase expires or resources end, your test account will be converted to a production account. Saved payment details will be used for billing purposes.

What happens when the trial period ends?

If you decide to continue using our services after the end of the test, the infrastructure you set up will remain in place; only the status of the account will change. User data saved during registration will be used for billing purposes. Henceforward, you will receive a detailed invoice for resources used and costs incurred. If you choose not to continue using it, delete your allocated resources before the end of the test and notify your account manager. If the manager arranges the deletion of your resources, then the deleted data cannot be recovered.

If I did not have the time to test, can I renew my test account again?

You have a test period of 30 days, and if you must extend your test period, contact your IONOS CLOUD account manager to find a solution.

Can I also test a private cloud for free for 30 days?

Yes, contact IONOS CLOUD sales department directly. This VMware-based solution on dedicated hardware will be built and delivered to you personally. Your account manager will first discuss your requirements and specifications, and then your infrastructure will be available for you in a few days.

What happens to my data if I do not choose to continue using IONOS CLOUD?

If you choose not to continue using IONOS CLOUD, delete your allocated resources before the end of the test and notify your account manager. If the manager needs to arrange the removal of your resources, we provide a multi-step removal process. The deleted data cannot be recovered.

After legal and operational retention periods expire, IONOS CLOUD will delete personal data collected when creating an account.

5. LEGAL

5.1. Untitled

5.2. Untitled

5.3. Untitled