NEBUL NEOCLOUD

SERVICE SPECIFIC TERMS

This document provides essential insights into our NeoCloud Services, designed to elevate your business by providing access to the world's most private and powerful AI infrastructure.

Table of Contents

1. Context and scope of this agreement 3
2. Service Description 5
3. Cloud Infrastructure 6
4. Managed Services 9
5. Data Handling and Privacy 12
6. Modifications to the Service Terms 13

Schedules

Schedule 1: Terms and Definitions 14
Schedule 2: Version Control 15

DEDICATED IAAS PRODUCT TERMS

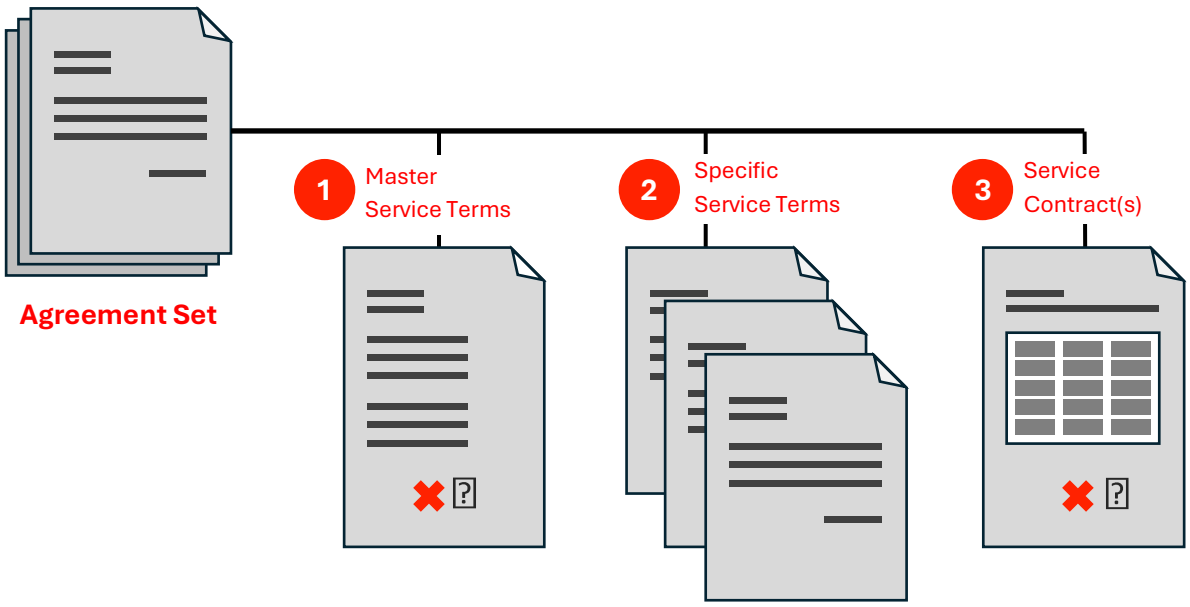
These Service Specific Terms apply when the customer purchases Nebul Dedicated IaaS services as identified on the customer's Service Contract. This document outlines the specific terms and conditions for Nebul's Dedicated IaaS offerings.

1. Context and scope of this agreement

This Agreement outlines the terms and conditions for using Nebul's services. The following sections will provide an overview of the contract structure, including supporting frameworks, terms, and agreements such as the Service Availability, Security Framework, and Support Plan.

1.1 Contract Structure.

The figure below illustrates the contract structure and how it integrates various supporting frameworks and agreements.



1.2 Relationship and Hierarchy of Agreements and Contracts

This document outlines the specific terms and conditions applicable to the Dedicated IaaS solution provided by Nebul.

(A) These terms are in addition to the Customer Agreement.

(B) In case of conflict, the documents prevail in descending order as described below:

i. Service Contract
ii. Acceptable Use Policy
iii. Specific Service Terms
iv. Service Level Agreement
v. Customer Agreement

This structured approach ensures that our agreements are clear, consistent, and enforceable, providing a reliable framework for delivering high-quality cloud services and an excellent customer experience.

2. Service Description

2.1 Cloud Management

Portal

The Nebul Cloud portal (https://portal.nebul.io) is a unified interface that empowers users to deploy and manage their cloud applications and services efficiently. The portal streamlines the entire application lifecycle through an intuitive, self-service approach.

Key Features

Centralized Management

The portal provides a comprehensive service catalog, workflow automation, and detailed documentation, reducing cognitive load and simplifying resource access.

Automated Operations:

Complex workflows are orchestrated through built-in automation, ensuring smooth operations with minimal manual intervention. Users can deploy applications rapidly while maintaining consistent performance and security standards.

Enhanced User Experience

- Intuitive navigation and interactive interfaces enable users of all skill levels to manage their resources effectively.
- Advanced search functionality increases the discoverability of services and documentation.
- Customizable dashboards provide relevant information and metrics at a glance.

Operational Benefits

The portal significantly improves efficiency by centralizing tools and providing self-service capabilities. This results in faster deployment times, reduced manual processes and enhanced productivity through automated task management. Teams can quickly access necessary resources, track compliance, and manage their applications through a single, unified interface.

2.2 API

The API gateway is an interface for managing cloud infrastructure and services through standardized REST endpoints.

Gateway Features

- Centralized authentication and authorization management
- Usage quota enforcement and rate limiting
- Request and response transformation based on policies
- Comprehensive logging and analytics for monitoring

Management Functions

- Self-service access for automated deployments
- Resource provisioning and configuration across environments
- Full lifecycle management of cloud services

Integration Benefits

- Automation Capabilities
- Standardized APIs for consistent resource management
- Programmatic control over infrastructure deployment
- Workflow automation for complex operations
- Integration with existing DevOps tools and processes

3. Cloud Infrastructure

3.1 Landingzone and Network

Secure Networks

Complete tenant isolation through customer-dedicated virtual private networks protected by Fortinet Next Generation Firewalls (NGFW) equipped with advanced threat protection and integrated security features, including IPS, antivirus, and web filtering.
Despite the sophisticated security measures, the end-user experience remains straightforward and intuitive. Users connect to their required resources through a simplified authentication process that works consistently across different access methods and locations. This approach significantly reduces user friction while maintaining robust security controls.

Advanced Security Architecture

The security infrastructure uses accelerated encryption to deliver high-performance, secure communications while maintaining stringent security standards. Unlike traditional perimeter-based security, this approach implements continuous authentication and authorization throughout each session, ensuring that trust is never assumed and always verified.

Modern Access Control

The platform supports traditional VPN technologies (IPSec and SSL) and Zero Trust Network Access capabilities, allowing organizations to implement security policies that best match their specific requirements. This flexibility ensures that organizations maintain strong security while providing users seamless access to necessary resources.

ZTNA allows organizations to transition smoothly from traditional VPN access to a more sophisticated, application-specific access control model. Users experience consistent access regardless of location, while security teams maintain granular control over resource access.

Management Interface

The unified management interface simplifies security administration by providing comprehensive visibility across all network edges. Security policies can be centrally managed and automatically enforced across the entire infrastructure.

3.2 Compute and SuperCompute

Nebul Cloud delivers Accelerated Cloud Compute and SuperCompute capabilities through a flexible, on-demand infrastructure designed for high-performance workloads. This infrastructure is ideal for running AI/ML workloads and accelerating everything else.

Accelerated Compute

Our compute service provides scalable virtual and bare-metal machines optimized for various workloads. Users can select from various pre-configured instances or customize their compute resources to match specific requirements.

Super Compute

Advanced GPU-based Supercomputers powered by NVIDIA technology enable accelerated computing for AI, machine learning, and complex simulations. The infrastructure supports various GPU-optimized containers and virtual machines that utilize CUDA, Tensor, and RT cores for specialized processing tasks.

Performance Benefits

- Enabling parallel computing and distributed workloads
- Reduction in processing time from days to minutes for complex workloads
- Hardware-accelerated computing for data-intensive applications

Workload Optimization

The platform automatically optimizes resource allocation based on workload requirements, ensuring efficient utilization and cost-effectiveness. Users can leverage built-in monitoring and scaling capabilities to adjust compute resources as needed, supporting vertical and horizontal scaling.

3.3 Storage

The platform provides a unified approach to data management while maintaining enterprise-grade security and performance capabilities across all access methods and workloads

Unified Storage Protocols

NetApp ONTAP provides seamless multi-protocol access, allowing simultaneous data access through NFS, SMB, and S3 protocols. This unified approach enables organizations to:

- Access the same datasets through different protocols without data reformatting
- Support both traditional NAS applications and modern S3-based analytics workloads
- Maintain consistent security policies across all access methods

Performance Architecture

Nebul uses the latest NetApp performance storage systems, with the AFF A90 system being the workhorse for NeoCloud Services, delivering exceptional performance for the most demanding workloads. Below is a summary of the storage platform features:

- End-to-end memory-based storage (NVMe) supports the latest NVMe/FC and NVMe/TCP protocols.
- Demonstrated performance of > 600 GiB/s in a Nebul cluster configuration.
- Scaling capability in a single namespace to over 20PB with over 400 billion files.
- Optimized for AI/ML workloads with NVIDIA GPUDirect Storage (GDS) support.

Data Protection

Autonomous Ransomware Protection using machine learning for threat detection
Snapshot technology for point-in-time recovery. Snap Lock for immutable backup copies.

Access Security

Zero Trust architecture follows principles of explicit verification and least privilege access.
Role-based Access Control (RBAC) for granular administrative access.

Encryption

Data encryption at rest using both software and hardware-based solutions.
Support for FIPS 140-2 validated cryptographic operations.

4. Managed Services

4.1 Managed Identity

Nebul uses KeyCloak as an identity and access management service to provide secure, centralized authentication and authorization services. This identity provider is a central hub for managing user access across all Nebul services.

Key Features

Single Sign-On

KeyCloak enables seamless Single Sign-On (SSO) functionality, allowing users to access multiple applications with a single set of credentials. The platform supports industry-standard protocols, including SAML 2.0 and OpenID Connect, ensuring secure communication between services.

Identity Integration

KeyCloak facilitates connections with external identity providers as an identity broker, enabling organizations to leverage their existing identity infrastructure. Customers can integrate their corporate identity systems, such as Active Directory or LDAP, while maintaining centralized control over access policies.

Security Features

The platform implements comprehensive security measures, including two-factor authentication, strong password policies, and advanced threat detection. User sessions are managed and monitored, with detailed audit logging for security compliance.

Access Management

Administrators can centrally manage all aspects of user access through an intuitive admin console, including user federation, role-based access control, and fine-grained authorization policies. The platform supports various access control mechanisms, from simple role-based access to sophisticated attribute-based and context-based controls.

4.2 Managed Private AI Factory

Nebul Private AI Factory is a comprehensive managed service designed for organizations that require secure, scalable AI infrastructure deployment under their own control.
This solution enables enterprises to build, train, and deploy AI models while maintaining complete data sovereignty and security controls.

Key Features

AI Development

- Data collection and preprocessing
- Test and refine AI models; Model selection and validation
- Experiment tracking and monitoring
- Validate hypotheses for business impacts
- Build automated data-pipelines
- Fine-tune foundational models

MLOps

The AI factory provides comprehensive MLOps capabilities, including automated model training pipelines, version control for models and datasets, and experiment tracking. This systematized approach ensures consistent development practices and enables rapid, agile development of AI systems.

Framework Support

Built-in support for popular AI frameworks and tools:

- Integration with popular AI frameworks like PyTorch and TensorFlow
- Large Language Model training and fine-tuning capabilities
- Distributed training architectures
- Custom model serving solutions

Data Protection

The Private AI Factory operates entirely within the customer's security perimeter, ensuring sensitive data and AI models never leave the controlled environment. All components adhere to European data protection standards and regulations.

Access Control

The Private AI Factory integrates comprehensive Role-Based Access Control (RBAC) mechanisms to enforce strict access controls and authentication protocols. Through identity management, the system assigns permissions based on specific job roles, ensuring users have minimal access to perform their duties. This security framework is enhanced by mandatory multi-factor authentication (MFA/2FA) for all account access, while comprehensive logging and monitoring capabilities track all system interactions.

4.3 Managed Kubernetes

Nebul K0rdent is a managed Kubernetes service that simplifies running Kubernetes in cloud and on-premises environments. As a certified Kubernetes-conformant platform, K0rdent ensures compatibility with existing applications designed for upstream Kubernetes.

Key Features

Cloud-Native Platform Building

K0rdent allows the creation of standardized application management environments by combining Cloud Native Computing Foundation (CNCF) technology with enterprise-grade services.

Running LLMs

The service supports deploying secure, scalable, high-performing Large Language Models (LLMs) for generative AI applications, leveraging high-performance GPU infrastructure.

Automated Management

K0rdent automatically handles the availability and scalability of Kubernetes control plane nodes, which are responsible for critical tasks such as container scheduling, application availability management, and cluster data storage.

Flexible Compute Options

The service supports running Kubernetes applications on various compute instances, virtual and bare-metal, including configurations with B200, H200, L40S, or L4 GPUs in a single instance or clusters scaling to thousands of GPUs.

Data Solutions

Build scalable, high-performing, and cost-efficient data platforms on Nebul K0rdent using managed services or open-source tools, all while maintaining European data sovereignty. By leveraging Nebul K0rdent, organizations can benefit from the robustness of Kubernetes while taking advantage of our extensive European infrastructure and service ecosystem.

Infrastructure Integration

Nebul K0rdent leverages enterprise-grade infrastructure to offer:

- High performance
- Scalability
- Reliability
- Availability
- Security and Network Integration

Flexible Deployment Options

K0rdent enables unified application management across cloud and on-premises environments, facilitating modernization initiatives.

5. Data Handling and Privacy

As an IaaS-provider, our role in data processing is limited to providing the underlying infrastructure on which customers deploy and manage their own applications and data. It's important to note that:

A) We do not have visibility into the specific categories of data our customers choose to process using our services;
B) Customers have full control and responsibility over what data they process, where they process it, and for what purposes; and
C) Our services offer standardized options, allowing customers to select the configuration best suited to their activities and compliance requirements.

Service-Specific Privacy Considerations

When using our Dedicated IaaS services, customers should be aware of the following privacy considerations:

- **Data at Rest:** Our storage solution provides data at rest encryption. Customers can use their own encryption keys, giving them full control over their data security.
- **Data in Transit:** Encryption for data in transit is possible using NFS 4.0 (Lower versions do not support encryption) or S3 protocols. However, the customer's responsible for ensuring the correct protocols are used to secure data during transfer.
- **Dedicated Infrastructure:** Our Dedicated IaaS infrastructure does not use multi-tenancy access. Each customer's environment is isolated, providing additional security and privacy.
- **Data Lineage and Provenance:** Customers should maintain their own records of data origins and transformations as we do not have visibility into these aspects.
- **Data Remanence:** Customers should be aware of and manage any data remnants after deletion or termination of services.

Reference to DPA

Our data processing activities are governed by a DPA that is compliant with GDPR requirements. The DPA outlines:

- The scope of our data processing activities on behalf of the customer.
- The categories of processing carried out (referenced to the services we provide).
- Our commitment to implementing appropriate technical and organizational security measures.
- Procedures for assisting customers with data subject requests pursuant to the GDPR.

Please refer to our DPA for the full details of our data protection commitments and compliance with the GDPR. Customers must review this document to understand the division of responsibilities between our IaaS service and their own obligations as data controllers under the GDPR.

Customers are reminded that they retain primary responsibility for ensuring their use of our Dedicated IaaS complies with applicable data protection laws, including the GDPR. We provide the tools and features to help customers meet their obligations, but the appropriate configuration and use of these features remain the customer's sole responsibility.

6. Modifications to the Service Terms

We may modify these Service Specific Terms for our Dedicated IaaS services from time to time in accordance with Section 11 of the Customer Agreement.

Schedule 1: Terms and Definitions

Agreement:	The Customer Agreement together with any other agreement governing using Nebul services, including this document.
Availability Zone:	An individual deployment of Nebul IaaS in a single location (zone) within a region.
Nebul NeoCloud	Nebul's Cloud offering provides exclusive access to private, virtual and bare-metal computing resources.
DPA:	Data Processing Addendum, outlining data processing activities and GDPR compliance.
GDPR	Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data.
GPU:	Graphics Processing Unit, used for high-performance computing and AI workloads.
HPC:	High-Performance Computing.
Nebul Site	www.nebul.com
NVIDIA SuperPOD:	The reference architecture on which Nebul IaaS solutions are built.
Region:	A deployment of Nebul IaaS in multiple distinct locations (zones), engineered to isolate it from failures in other locations.
RoCE:	RDMA over Converged Ethernet is a network protocol that allows remote direct memory access over an Ethernet network.
SLA:	Service Level Agreement, outlining specific commitments and performance metrics.

Schedule 2: Version Control

Version	Date	Status	Information
1.0	May 2023	Archive	Initial Terms
1.1	July 2024	Archive	Various updates
1.2	March 2025	Current	Layout update (Nebul rebrand)