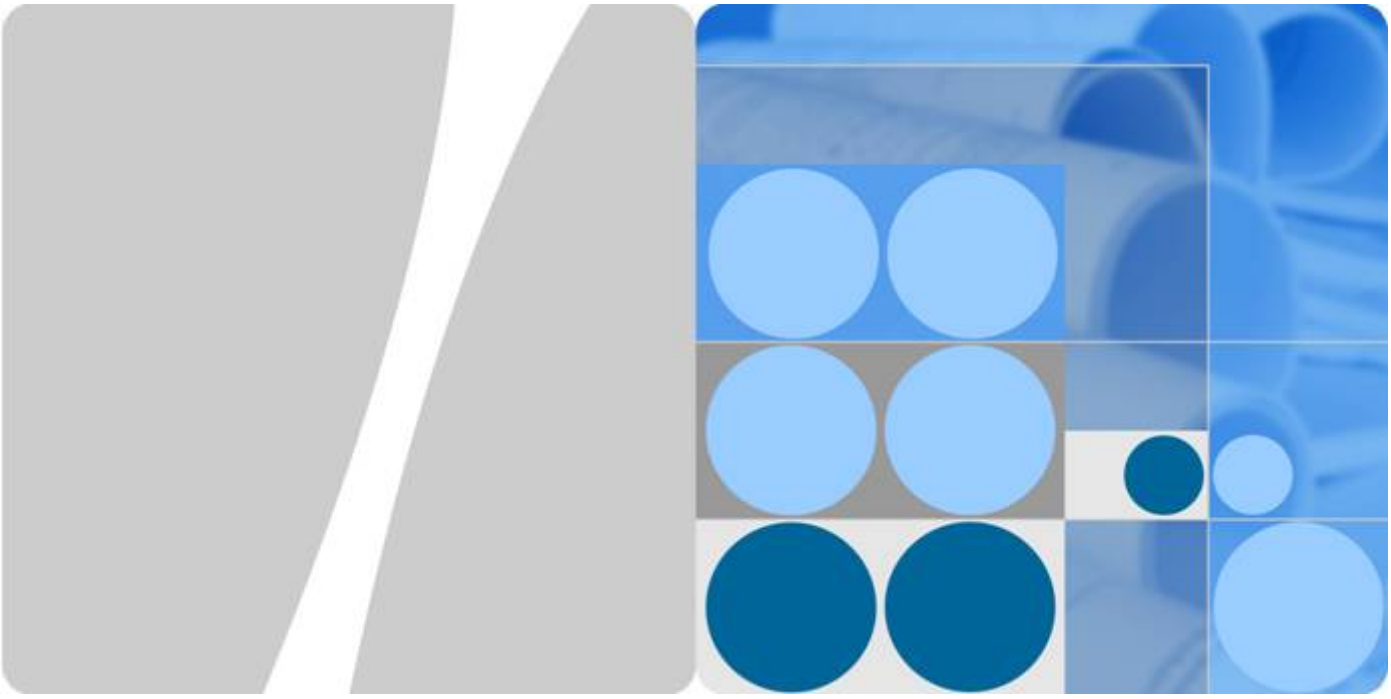

Cloud Data Migration

Security White Paper

Issue 01

Date 2018-08-03

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions


and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Contents

1 CDM Security Overview

Introduction to CDM

Cloud Data Migration (CDM) implements data mobility by enabling batch data migration among homogeneous and heterogeneous data sources. It supports on-premises and public-cloud-based data sources, including file systems, relational databases, data warehouses, NoSQL, big data, and object storage.

CDM is suitable for use in the following scenarios:

1. Data migration to the cloud: When using HUAWEI CLOUD services, tenants can migrate their historical data or incremental data from the private cloud, on-premises data center, or third-party public cloud to HUAWEI CLOUD.
2. Data exchange between cloud services: Tenants can migrate data among the big data service, database service, and object storage service of HUAWEI CLOUD. For example, data processed by MapReduce Service (MRS) can be imported to Data Warehouse Service (DWS) for interactive analysis and report statistics collection.
3. Data migration from the cloud to on-premises: Computing resources on the public cloud can be used to process huge data sets. The processing results are then returned to on-premises business systems (relational databases and file systems).

Migration Principles

When a tenant uses CDM, the CDM system provisions a fully-managed CDM instance in the tenant's VPC. The instance allows only console and RESTful API access. Therefore the tenant cannot access the instance through other interfaces (such as SSH). This ensures data isolation between CDM tenants, prevents data leakage, and ensures transmission security during data migration between different HUAWEI CLOUD services in a VPC. Tenants can also use the VPN to migrate data from the on-premises data center to HUAWEI CLOUD services to ensure migration security.

CDM works in push-pull mode. CDM pulls data from the migration source and pushes the data to the migration destination. Data access operations are initiated by CDM. SSL will be used if the data source (such as RDS) supports it. During the migration, the usernames and passwords of the migration source and destination are required. Such information is stored in the database of the CDM instance. Protecting such information is critical to ensure CDM security.

Figure 1-1 Migration principles

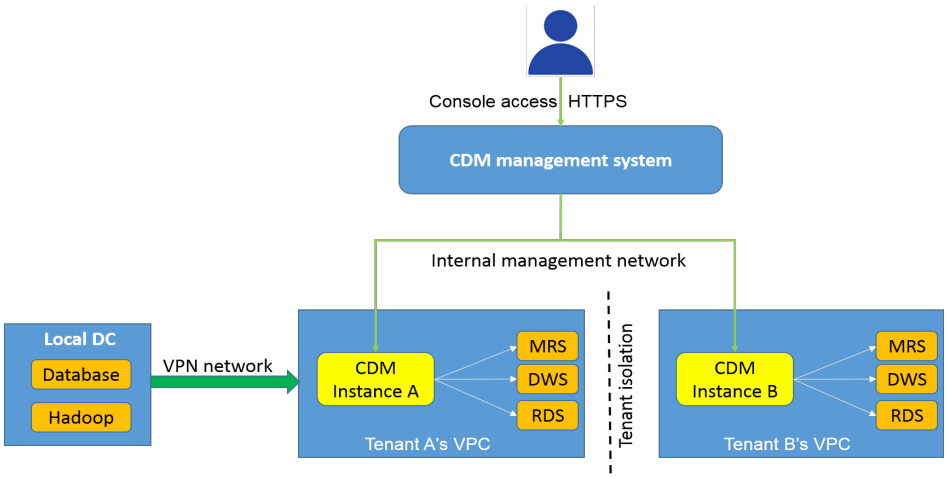

Security Boundary and Risk Mitigation

Figure 1-2 Risk mitigation

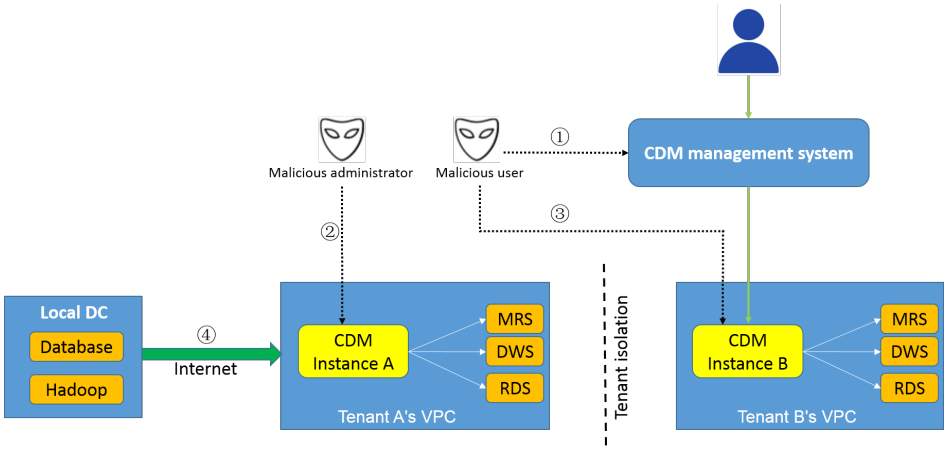

As shown in the preceding figure, CDM may have the following threats:

1. Threats from the Internet: Malicious tenants may attack CDM through the CDM console.
2. Threats from the data center: Malicious CDM administrators obtain tenants' data source access information (usernames and passwords).
3. Threats from malicious tenants: Malicious tenants steal data from other tenants.
4. Data exposure to the public network: Data is exposed when it is migrated from the public network.

CDM offers the following mechanisms to prevent security risks:

1. Threats from the Internet: Tenants cannot log in to the CDM console through the public network. CDM provides a two-layer security mechanism.
   a. On the one hand, the HUAWEI CLOUD console framework requires user authentication when tenants access any console.
   b. On the other hand, Web Application Firewall (WAF) filters requests from all consoles and stops request attack code or content.

2. Threats from the data center: Tenants must provide the usernames and passwords of the migration source and destination to complete data migration. To prevent the CDM administrators from obtaining such information and attacking important data sources of tenants, CDM provides a three-level protection mechanism.
   a. CDM stores passwords encrypted by AES-256 in the database of the instance to ensure tenant isolation. The database is run by user **Ruby** and listens to only 127.0.0.1. Therefore, tenants cannot remotely access the database.
   b. After the instance is provisioned, CDM changes the passwords of users **root** and **Ruby** to random passwords and does not store them in any place. This prevents the CDM administrators from accessing tenants' instances and databases containing password information.
   c. CDM instances work in push-pull mode. Therefore, the instances do not have any listening port enabled in the VPC, and tenants cannot access the local database or operating system from the VPC.
3. Threats from malicious tenants: CDM runs instances on independent VMs, so that tenants' instances are completely isolated and secure. Malicious tenants cannot access instances of other tenants.
4. Data exposure to the public network: In push-pull mode, even if elastic IP addresses (EIPs) are bound to the CDM clusters, no port is enabled for the EIPs. In this way, attackers cannot access and attack CDM using the EIPs. However, when data is migrated from the public network, tenants' data sources are exposed to the public network and threatened by third-party attacks. Therefore, tenants are advised to use ACLs or firewalls on the data source server for security. In this case, for example, only the access requests from the EIPs bound to the CDM clusters are allowed.

2 CDM Security Conclusion

Access Control

Only tenants authorized by Identity and Access Management (IAM) can access the CDM console and APIs. In push-pull mode, CDM does not have any listening port enabled in the VPC. For that reason, tenants cannot access instances from the VPC.

Data Transmission Security

CDM runs in tenants' VPCs to ensure data transmission security in terms of network isolation. Data sources that support SSL, such as RDS and SFTP, can be accessed in SSL mode. CDM also allows data of public data sources to be migrated to the cloud. Tenants can use the VPN and SSL to prevent transmission security risks.

Tenant and Network Isolation

CDM instances run in independent VPCs. VPC allows tenants to configure VPC inbound IP ranges to control the IP address segments for accessing CDM. After a CDM instance is deployed in a tenant's VPC, the tenant can configure the subnet and security group to isolate the CDM instance, thereby improving the security of the CDM instance.

Data Encryption

The access information (usernames and passwords) of tenants' data sources is stored in the database of the CDM instance and encrypted using AES-256. The CDM administrators cannot access the database.

Data Deletion

When a tenant delete a CDM instance, all data stored in it will be deleted. Nobody can view or restore the deleted data.

