

IBM Cloud Object Storage System™
Version 3.14.1

Container Mode Guide



This edition applies to IBM Cloud Object Storage System and is valid until replaced by new editions.

© **Copyright IBM Corporation 2016, 2018.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Container Mode Guide	1
Overview	1
Vault mode	2
Container mode	3
Storage accounts	3
Accounts and conversion to container mode	3
Container mode workflow	4
Enabling Container Mode with a new system	4

Enabling container mode with an existing system	9
Operations	12
Manager operation	12

Notices	13
Homologation statement	15
Trademarks	15

Container Mode Guide

The system supports two different modes of operations: vault mode and container mode. By default, the system operates in vault mode. In vault mode, all object I/O occurs at the vault level. Systems can contain a limited number of vaults. Container mode can be enabled when more buckets than the vault limit would otherwise allow are required. In container mode, containers are created inside of vaults. Object I/O occurs on containers instead of vaults.

Table 1. Advantages of container mode

Container mode	Vault mode
Support millions of buckets	Support for a limited number of buckets
Support for millions of users	Support for thousands of users
Support for fast and high performance listing with index updates (system OPs greatly reduced for listing)	Slow listing and unreliable performance
Support for usage based billing based on top-edge requests via access logs	Billing is supported based on estimated usage for storage and network resources on a vault basis
Support for optional storage metrics cluster (SMC) with REST API support for billing	No support for billing via storage metrics cluster
Support for self service portal for millions of end users via service API	Provides portal for end users to manage buckets (vaults)
Support for S3 ACLs	S3 ACLs are supported, but are not fully compliant to S3
Support for additional S3 APIs (for example, GET SERVICE)	Support for standard S3 APIs previously supported

Overview

In container mode, each container vault can contain over 1 million containers. When enabled, I/O occurs at the container level instead of the vault level. ACLs are enforced at the container level. Metrics are tracked at the container, storage account, and container vault levels. Containers and storage accounts are not visible on the Manager. Accesser[®] Devices can be used to create, update, and delete containers.

Table 2. Limitations with container mode

Container mode	Vault mode
No support for Swift, SOH, DDN/WOS, or HDFS.	Top-edge support for Swift, SOH, DDN/WOS, and HDFS.
No support for S3 versioning.	S3 versioning support.
No support for data migration and proxy.	Data migration and proxy support.
No support for Keystone and AD for user accounts and credentials management.	Keystone and AD support for user accounts and credentials management.
No Manager visibility for user accounts and credentials.	Manager visibility for all user accounts and credentials.
No support for mirroring.	Vault mirroring support.
No support for IP allow/disallow at the container level (support exists at the container vault level).	IP allow/disallow support.
No support for locked or private vaults.	Private or locked vaults support.
No support for S3 bucket tagging.	S3 bucket tagging support.

Table 2. Limitations with container mode (continued)

Container mode	Vault mode
No quota enforcement at the container level (support exists at the container vault level).	Quota enforcement at the vault level.
No support for Accesser application.	Accesser application supported.
No support for Embedded Accesser.	Embedded Accesser supported.

The following requirement must be met before you enable container mode:

- Native File Interface cannot be enabled on the same system.

When in container mode, the system has the following restrictions:

- The only API type that is supported is Cloud Object Storage.
- Name index must be enabled on all container vaults. It is done automatically when container vaults are created.
- Versioning is not supported.
- Private vaults are not supported.
- Mirrors are not supported.
- Delete restricted vaults are not supported.
- Container mode can be disabled only when all container vaults are deleted.
- Container vault deletion is only supported by assistance from IBM Support.
- The service vault can be deleted only after container mode is disabled.
- Use of the embedded Accessor service feature is not recommended.

Vault mode

Cloud Object Storage accounts are defined in the Cloud Object Storage Manager and are used for vault level I/O.

Storage I/O against a vault with these accounts can be done in two ways:

- Basic authentication
- Access key authentication

The general roles that can be assigned to an account are:

- Super user
- System administrator
- Security officer
- Operator
- Service account

Except for the *Service account* role, each of these roles affects what the account can do and see in the Cloud Object Storage Manager, irrespective of whether the system is in Container Mode or not. However, the *Service account* role grants the account access to the Service API in Container Mode and cannot be set unless the system is in Container Mode.

The types of vault access are:

- Owner
- Read/write
- Read

- No access

If an account has access to a vault that is a part of a conversion to container mode, a *Storage account* is created for it (see “Storage accounts”). In Container Mode, all I/O is performed by using access keys.

Note: In container mode, the *Service account* role is used for accessing the Service API.

Container mode

Storage accounts

With the introduction of Container Mode, the concept of a storage account was introduced. A storage account is an entity on which users and containers are created. Storage accounts are often the entity that is billed for storage usage.

Each storage account in the system has at least one set of access keys, which are needed for all Cloud Storage Object requests on the container.

Accounts and conversion to container mode

The following three tables show a scenario in which an account has access keys, permission to perform I/O on a vault, and the *Service account* role in the Cloud Object Storage manager.

Table 3. Account in Cloud Object Storage Manager

Name	Access key ID	Secret access key
acct1	GQn7t4S0u3LXrcCIRpKn	ShZAq0dEdrAX6D6B4iJUK

Table 4. Standard vault permissions in Cloud Object Storage Manager

Vault name	Account name	Permission
vault1	acct1	readWrite
vault2	acct1	readWrite

Table 5. General Cloud Object Storage Manager roles for accounts

Account name	General roles (list)
acct1	system administrator, service account

When you perform a Container Mode conversion in this scenario, when **vault1** is converted to Container Mode, a *Storage account* is created for **acct1**. The access key credentials in the first table can be used to perform I/O at the container level.

Having a *Storage account* that is created for the account does not prevent the account from having access to the Service API. Access to the Service API is solely determined by the existence (or a lack thereof) of the *Service account* role. Thus, in this situation, the *Account* has access to the service API and the *Storage account* has access to container level I/O, both with the same set of access keys defined in Table 3.

In addition, even after **vault1** is converted to a container vault, a standard vault **vault2** still presides on the system. Since **acct1** has permission to the vault, it can perform vault level I/O to vault **vault2** by using basic authentication or by using the access keys that are defined in Table 3.

Access key credentials that are used by the *Storage account* for container level I/O can be changed. When new access keys are added or removed by using the Service API, the Cloud Object Storage Manager does

not show the updates. The access keys that are used by the *Storage account* are not managed by the same entity as the ones used by the *Account* in the Cloud Object Storage Manager.

Similarly, if access keys are added to the account in the Cloud Object Storage Manager, these keys can be used for vault-level I/O or the Service API, but they cannot be used for container-level I/O.

For these reasons and to avoid confusion, it is not recommended that you use the same base account (**acct1** in this case) and its storage account for accessing the Service API and performing container I/O. It is recommended that you create new accounts in the Cloud Object Storage Manager for accessing the Service API. In most cases, confusion can be avoided when all vaults that a single account has access to are converted at the same time (for example, **vault1** and **vault2**). It avoids a mixed Vault Mode/Container Mode scenario.

Container mode workflow

Enabling Container Mode with a new system

About this task

Perform the following steps to enable and use Container Mode:

Procedure

1. Create a service vault.
2. Enable Container Mode.
3. Create a service account.
4. Create a container vault.
5. Configure the service API.
6. Create a Storage Account.
7. Create access keys.
8. Create a container.

Creating a service vault

About this task

The service vault is used to store metadata that is related to containers, storage accounts, and access keys. Each IBM Cloud Object Storage System™ has a single service vault. It is important to select an IDA that ensures no data is lost from the service vault as data loss would lead to the loss of containers. Also, if the service vault is not accessible, all container I/O fails. However, selecting an IDA that is too wide results in reduced container I/O performance.

Note: The service vault must be created on a storage pool with packed storage enabled.

Procedure

1. Click the **Configure** tab.
2. Click **Configure** in the **Configure Container Mode** section.
3. Enter the information for a new service vault and click **Create Service Vault**.

Enabling container mode

About this task

During this step, the Manager validates that the current configuration is compatible with container mode. Detected errors are displayed and must be fixed before you can continue.

You must abide by the following validation rules:

1. No standard vaults can exist.
2. The API type of all access pools must be cloud storage object.
3. A service vault must be present.

After creating the service vault, new container vaults can be created.

Procedure

Select the **Create only container vaults** option on the **Configure container mode** page. All new vaults will be container vaults.

Note: To add new standard vaults to the system in addition to container vaults, you will be operating in mixed mode. More information on mixed-mode operation is below.

Creating a service account

About this task

Manager accounts must be used for authentication or authorization of service API requests and are needed to have the Service Account role. Complete the following steps to add the Service Account role to a manager account.

Procedure

1. Click the **Security** tab.
2. Create an account or modify an existing account in the **Accounts** section.
 - For a new account, click **Create Account**.
 - For an existing account, click the user name to whom you want to assign the role and then click **Change**.
3. Enable the **Service Account** role in the **Roles** section.

Creating a container vault

Create a container vault to house containers.

Procedure

1. Click the **Configure** tab.
2. Click **Create Vault** in the **Summary** section.
3. Choose the method by which you want to create a container vault:
 - **Using Template:** Select a template and click **Continue**.
 - **Custom Vault From Storage Pool:** Select a storage pool and click **Continue**.

In the Select vault type dialog box, select **Container**.

Optionally, select **Create only container vaults**. New vaults will be container vaults. If standard vaults are required, disable this option. This option will have no effect for vaults created on incompatible storage pools or vault templates.

Click **Continue**.

4. In the **General** section, complete the following fields:

Tip: The provisioning code is located in the **General** section, and is used to specify the container vault in which you want to create containers.

Field Label	Acceptable Field Value
Name	Each vault must be uniquely named (maximum of 255 characters); this name is used by the Manager for all references to this vault. Vault names can include underscores and alphanumeric characters. The vault name can also contain periods (.), but the name cannot start or end with a period or contain more than one period in a row. The first character of the name must be a letter, underscore, or number.
Description	An optional free-form description can also be entered. Information that you might include in the description field might be initiator host name and IP address, names and phone numbers of administrators, and key users of the vault.
Tags	Tags can be created and or assigned to a vault before the vault is created. For more information, see the Tags section.
Provisioning Code	During container creation, specify a unique provisioning code with a maximum of 255 characters to indicate in which container vault the container should be created. The default value is the vault name, once entered. The value is usually the locationConstraint or region . If the provisioning code is not set, containers can be created in this new container vault if it is configured as the default container vault for an access pool.
Region	Optional. A region can be provided with a maximum of 255 characters to indicate where the contents of this vault reside. The locationConstraint parameter, which is shown for containers that are associated with this vault in the S3 GET Service Extended and S3 GET Bucket Location APIs, are populated with the value set for region.
Storage Class	Optional. A storage class can be provided with a maximum of 255 characters to assign a classification to all objects stored within this vault. The header x-amz-storage-class shown in the S3 GET/HEAD Object and the StorageClass in the response body of the S3 GET Bucket are populated with the value set for storage class.

5. In the **Configuration** section, several options display.

Field Label	Acceptable Field Value
Width	<p>This setting is referred to as the width of the vault and corresponds to the number of slices into which all data in the vault is split.</p> <p>Vault width must be a factor of the storage pool width. The Manager Web Interface allows any vault width greater than or equal to 6 and less than or equal to 60.</p>
Threshold	<p>The minimum number of slices that must be available to perform a read. Pre-defined, supported thresholds are presented when the drop-down list is clicked. The vault threshold, always less than the width, determines the reliability of the vault. If the set of available Slicestor[®] devices is such that the number of slices falls below this threshold, the vault content cannot be read, and the vault appears as red in the Monitor application.</p> <p>The Manager Web Interface allows any value between 1 and Vault Width, inclusive.</p> <p>If the vault is on a storage pool that spans multiple sites, the Manager Web Interface warns the user if the selected threshold is high enough such that a single site outage affects read and write availability.</p>

Field Label	Acceptable Field Value
Write Threshold	<p>The Manager Web Interface allows any value such that all the following are true:</p> <ul style="list-style-type: none"> • Write Threshold > Threshold. <p>CAUTION: Write Threshold = Threshold is allowed if Threshold = Vault Width or if Vault Width < 6.</p> <ul style="list-style-type: none"> • Write Threshold ≤ Vault Width. • (Write Threshold + Threshold) > Vault Width. <p>Write Threshold defaults to Threshold + 2, if that is within the allowed range. Otherwise, the selected Write Threshold is the halfway point between the minimum allowed Write Threshold and Vault Width, rounded up. This value is selected by default in the Write Threshold drop-down when Threshold is selected. This value is also used as the Write Threshold when a vault is created through the Manager REST API and a Write Threshold is not specified.</p> <p>If the vault is on a storage pool that spans multiple sites, the Manager Web Interface warns the user if the selected write threshold is high enough such that a single site outage affects write availability.</p>
Alert Level	<p>Optional. If the set of available Slicestor devices is such that the number of slices is between the write threshold and the alert level exclusive, the vault icon is yellow in the Monitor application. In this case, the vault is still fully functional.</p>

Attention: If the **Threshold** is set such that the loss of one site would make Vaults either unusable or read-only, the Manager Web Interface displays a confirmation dialog box that asks the operator if they accept the settings with the risks they present.

If only the **Threshold** causes an issue:

Warning: This IDA configuration is susceptible to read availability issues during a single site outage. Do you still wish to continue?

If the **Threshold** and **Write Threshold** cause issues:

Warning: This IDA configuration is susceptible to read and write availability issues during a single site outage. Do you still wish to continue?

Click **Cancel** or **OK** to change or keep the settings.

6. If you enabled vault protection on the system, choose a Retention setting.

Note: This section displays only if Vault Protection Configuration is enabled in the **Configure** tab.

- **Disabled.** The container vault does not support Retention.
- **Enabled:** The container vault supports Retention for buckets in this container vault.

7. In the **Options** section, complete these fields:

Field Label	Description
Enable SecureSlice™ Technology.	<p>Optional. SecureSlice™ provides extra encryption benefits that are combined with dispersal. This box is checked by default for newly created vaults. This feature can be cleared, although it is not recommended. If it is cleared, a warning message appears, and a confirmation is needed before proceeding. When a vault is created, the SecureSlice™ option cannot be changed.</p>

Field Label	Description
Enable Server Side Encryption with Customer-Provided Keys	Optional. SSE-C is enforced on objects. Requests to read or write objects or their metadata using customer managed keys send the required encryption information as headers in the HTTP requests.
Enable Server Side Encryption with Key Protect (SSE-KP)	Optional. Use IBM Key Protect to create, add, and manage keys, which you can then associate with your instance of IBM® Cloud Object Storage to encrypt buckets.
Enable Archive Tiering (cannot be disabled)	Optional. This feature provides a service user the ability to create a bucket and assign an archive policy, ability to add an archive policy to an existing bucket, modify an archive policy for a bucket (only applied to newly written objects in the bucket), view the archive policy for a bucket, delete/disable archive policy for a bucket.

8. In the **Quotas** section, complete these fields:

Field	Value
Soft Quota	Optional. If wanted, select a value for a soft quota. A notification is sent to the Event Console , if the soft quota setting is exceeded. It does not cause restrictions to usage. Setting the quota higher than the total space available in one or more storage pools that are associated with this vault has no effect.
Hard Quota	Optional. If wanted, enter a hard quota value. The Accesser® device (or application) does not permit the user to exceed the hard quota value for this vault. A notification is also sent to the Event Console if the hard quota setting is exceeded. Setting the quota higher than the total space available in one or more storage pools that are associated with this vault has no effect.

9. Click **Save**. A new container vault is created.

Configuring the service API

You can create or edit an access pool to open the service API ports on at least one access pool.

About this task

By default, the service API ports are not opened because most access pools do not need to have these ports open.

Procedure

1. Click the **Configure** tab.
2. Perform one of the following steps to open the service API ports.
 - For a new access pool, click **Create Access Pool** in the **Summary** section. For more information, see *Create access pools* in the *Configuration* chapter of the *Manager Administration Guide*.
 - For an existing access pool, click **Open All** above the left navigation tree and then click an access pool. Click **Change**.
3. Enable the service API ports in the **General** section.
4. Click **Save**.

Create a storage account

Storage accounts are entities on which users and containers are created. Often, storage accounts are the entity that is billed for the storage usage. The creation of the storage account is the first step in allowing access to the system.

Storage accounts are created through Service API requests, as in the following example:

```
curl "http://<accesser ip>:8337/accounts/<storage account name>" -X PUT -u <user name>:<password>
```

- <accesser ip> is the IP address of an Accesser[®] Device that belongs to an access pool with the service API enabled. To use an IPv6 address, enclose the Accesser IP in square brackets: [<accesser ip>].
- <storage account name> is the name of the storage account that is created. The storage account name must be unique.
- <user name> is the user name of a manager account that has the Service Account role.
- <password> is the password of the <user name>.

Create access keys

After you create a storage account, you can create access keys for users that have access to the storage account. Access keys are needed for all Cloud Object Storage requests on the container.

Access keys are created through Service API requests, as in the following example:

```
curl "http://<accesser ip>:8337/credentials" -X POST -u admin:password
-d '{"credential":{"project_id": "<storage account name>","type":"ec2"}}' -H "Content-Type: application/json"
```

- <accesser ip> is the IP address of an Accesser[®] Device that belongs to an access pool with the service API enabled. To use an IPv6 address, enclose the Accesser IP in square brackets: [<accesser ip>].
- <storage account name> is the name of the storage account the access key is associated with.

As part of the response, the key ID, key, and secret key, secret, are returned. These credentials must be used by the user to create and modify containers.

Create a container

After you create access keys, you can create containers.

Use the access key to determine which storage account to create the container in, as access keys are associated with a single storage account. Containers are created with Cloud Object Storage **PUT** bucket requests, as shown in the "Create a new bucket" section of the *IBM Cloud Storage Object API Developer Guide*.

Note: In **PUT** bucket requests in container mode, the IP to use is the IP address of an Accesser[®] Device that belongs to an access pool with the service API enabled. To use an IPv6 address, enclose the Accesser IP in square brackets: [<accesser ip>].

The location constraint field should be the provisioning code of the container vault where the container is created. The provisioning code is configured on the create or edit vault page. This parameter is NOT required if a default container vault is configured for the Accesser[®] Device. The default container vault can be set on the create or edit access pool page.

Enabling container mode with an existing system

About this task

The conversion process is only needed if one or more standard vaults exist in the system and there's a desire for them to be container vaults. Before you enable container mode and initiating a vault conversion, it is highly recommended that you familiarize yourself with the differences in how the system operates in container mode. Moving from a system that uses vault level I/O to one that uses container level I/O is a significant change. It cannot be undone.

Standard vaults are not converted directly to container vaults. Instead, access pools are selected for conversion. When an access pool is selected, all of its vaults are converted from standard to container vaults. If a single vault is deployed to multiple access pools, it is not possible to convert only a subset of the access pools. All of them must be converted. An alternative option is to remove one or more of the access pools from the vault's deployments before you start a conversion.

During the conversion process, the Cloud Object Storage Manager does not allow configuration changes to be made.

I/O to existing vaults do not operate any differently when the conversion process is ongoing. If the vault is being converted to a container vault as a part of the process, any I/O initiated while the vault was a standard vault (before conversion, for example) succeeds normally, even if the vault becomes a container vault before the I/O is completed. After the vault is converted into a container vault, traditional vault level I/O will no longer be possible: container level I/O must be used.

Procedure

1. Create a service vault. The service vault is needed for container mode. It is used to store metadata that is related to containers, storage accounts, and access keys. It must be created, but needs to be created only one time.
2. Ensure that all vaults that are deployed to the selected access pool satisfy all conversion compatibility requirements.
3. Initiate conversion. During an ongoing vault to container conversion, the Cloud Object Storage Manager rejects all configuration requests. It is to ensure that the new containers have current metadata. The conversion can take several minutes.
4. Create a container for each vault. The new container's name matches the container vault's name as it is configured in the Cloud Object Storage Manager.
5. Create a Storage Accounts for each account with permissions to a vault. Credentials are created for the storage account by using the access keys that are owned by the account.
6. Wait for the conversion to complete. Configuration requests to the Cloud Object Storage Manager are permitted again.

What to do next

After an access pool is converted, each of its deployed vaults will be container vaults. New I/O at the vault level is no longer possible on the vaults. Objects must be accessed by using container I/O methods, which require the use of account access keys. To ensure a smooth conversion, you should start the system by using access keys for vault I/O before you convert to container mode. Only container vaults can be deployed to the access pool.

Storage account credentials

A storage account and its credentials are created for each account with permissions to the vault.

If multiple conversions take place over time and a storage account was created for an account during a previous conversion, the system will not attempt to re-create the storage account or credentials. Deleting access keys from the account in the Cloud Object Storage Manager has no effect on the access keys that are associated with the storage account.

Conversion compatibility requirements

Not all vaults can be converted. Several compatibility requirements exist that could prevent a vault, or an access pool, from being compatible for conversion.

System-wide requirements.

- A service vault must exist in the system.

Requirements for each vault.

- At least one account must have **readWrite** or **owner** permissions assigned.
- Name index must be enabled ¹.
- Vault proxy is not allowed.
- Recovery Listing is not allowed ¹.

- Locked vault is not allowed ¹.
- Cardinal vault is not allowed.
- Delete restrictions not allowed ¹.
- Versioning is not allowed ¹.
- Being one side of a mirror is not allowed.
- Notification Service configuration is not allowed.

Requirements for each vault's storage pool.

- Packed storage must be used ¹.

Requirements for each access pool.

- API type must be 'Cloud Object Storage' or 'OpenStack Object Storage'.
- Mirrors cannot be deployed.

Requirements for each account with permission to a selected vault.

- One or more access keys must be created.

If a system does not obey these requirements, several of them can be reconfigured in the Cloud Object Storage Manager. For instance, if a vault proxy is configured on a vault, "fixing" the vault for conversion requires the vault remove the proxy.

¹ No simple fix for this requirement exists. In these cases, no way to convert the standard vault to a container vault exists, at least without first migrating data to a new vault by using the vault migration function.

Note: If a vault contains objects that are written with Simple Object (SOH), those objects cannot be read if the vault is converted to a container vault.

Mixed operation

After the service vault is created, the Cloud Object Storage Manager changes in two distinct ways.

- Allows standard vaults to be converted to container vaults.
- Allows the creation of container vaults in the system.

New user-created vaults can be standard vaults or container vaults. The one exception to this is that a vault that is created on a storage pool that uses file storage or containing mirrors, must be a standard vault. In other cases, both types of vaults can be created. In these situations, when you create a new vault, the type of vault (standard or container) must be selected. The configuration options on the vault creation page differ if the vault is a container vault versus a standard vault. Standard vaults show these parameters; they cannot be set for container vaults:

- Name index.
- Recovery listing.
- Versioning.
- Delete restriction.
- Authorized users.

If you create a standard vault now, with the desire to convert it into a container vault in the future, take special care when you select any of these options. See the conversion compatibility requirements.

Accounts cannot be given direct access to a container vault in the Cloud Object Storage Manager. It is handled only at the container level. For more information, see the *Service API Manual*.

Operations

After container mode is enabled, the behavior of the system will change slightly. I/O is now run on containers instead of vaults. Also, some manager screens are slightly modified (refer to the following section for details).

Manager operation

After you migrate the system to container mode, you can still create, configure, and monitor vaults through the Manager.

Create or edit a vault

The **Create vault** and **Edit vault** pages look slightly different. Certain configuration parameters are no longer visible as they cannot be changed on container vaults.

The following configuration parameters are no longer visible:

- Name index.
- Recovery listing.
- Versioning.
- Delete restriction.
- Authorized users.

Also, a new configuration parameter is added.

Provisioning Code

Used to specify in which container vault containers should be created. During PUT bucket requests, the provisioning code of the wanted container vaults should be specified as part of the LocationConstraint. During container vault creation, the provisioning code defaults to the vault name. It can be changed by clicking the provisioning code text box and making the wanted edits.

Create or edit an access pool

The **Create access pool** and **Edit access pool** pages are also modified to include an extra parameter, default container vault.

If a default container vault is specified, the LocationConstraint in the PUT bucket requests becomes optional. If no LocationConstraint is specified, the container is created in the device's default container vault.

Also, the service ports can be opened or closed on the create access pool or edit access pool pages.

Configuring container mode

The **Configure container mode** page can be used to configure whether the container names must be DNS-compliant names. By default, this restriction is not enabled.

Also, container mode can be disabled on the **Configure container mode** page. However, container mode can be disabled only when all container vaults are deleted from the system. Container vault deletion requires assistance from IBM® Support.

Delete a vault

Container vaults cannot be deleted without assistance from IBM Customer Support.

The service vault cannot be deleted until container mode is disabled. Container mode can be disabled when all container vaults are deleted.

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Homologation statement

This product may not be certified in your country for connection by any means whatsoever to interfaces of public telecommunications networks. Further certification may be required by law prior to making any such connection. Contact an IBM representative or reseller for any questions.

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Accesser[®], Cleversafe[®], ClevOS[™], Dispersed Storage[®], dsNet[®], IBM Cloud Object Storage Accesser[®], IBM Cloud Object Storage Dedicated[™], IBM Cloud Object Storage Insight[™], IBM Cloud Object Storage Manager[™], IBM Cloud Object Storage Slicestor[®], IBM Cloud Object Storage Standard[™], IBM Cloud Object Storage System[™], IBM Cloud Object Storage Vault[™], SecureSlice[™], and Slicestor[®] are trademarks or registered trademarks of Cleversafe, an IBM Company and/or International Business Machines Corp.

Other product and service names might be trademarks of IBM or other companies.



Printed in USA