Support for Multiple Organizations

inSync Private Cloud 5.4

Druva Inc.

Email feedback to docs@druva.com

druva

Table of Contents

1. About this guide 5
2. Overview of support for multiple organizations 6
2.1 Introduction 6
2.2 The multi-org benefits 7
2.3 How multi-org works 8
3. Planning for a Private Cloud setup 10
4. What you will need to get started 13
4.1 Installing the Master server 13
4.2 Using the Quick Configuration Wizard 13
4.3 Logging on to Organization Portal 14
4.4 Modifying your Private Cloud setup 15
4.4.1 Modifying email setup 15
4.4.2 Modifying the network settings 16
5. Working with the Organization Portal 19
5.1 Adding and managing organizations 19
5.1.1 Adding an organization 19
5.1.2 Modifying an organization 21
5.1.3 Enabling or disabling organizations 21
5.1.4 Deleting an organization 22
5.1.5 Exporting user data 22
5.2 Installing additional storage nodes 23
5.2.1 Prerequisites for installing a storage node 23
5.2.2 Installing a storage node 24
5.2.3 Registering the storage node 25

5.2.4 Loading an SSL certificate on a storage node 26
5.2.5 Modifying a storage node 26
5.3 Creating and managing additional storage 29
5.3.1 Understanding types of storage 29
5.3.2 Creating a file store 32
5.3.3 Creating an object store 36
5.3.4 Modifying storage details 40
5.3.5 Managing data volumes on a storage 40
5.3.6 Deleting a storage 41
5.3.7 Compacting a storage 41
5.4 Creating storage pools 43
5.4.1 Understanding storage pools 43
5.4.2 How user migration works 44
5.4.3 Creating a storage pool 47
5.4.4 Modifying storage pool 48
5.5 Creating an HA policy 49
5.6 Setting up inSync Edge Server 51
5.6.1 Understanding the edge server 51
5.6.2 Understanding inSync Private Cloud architecture 52
5.6.3 How inSync Private cloud works 56
5.6.4 Installing the edge server 58
5.6.5 Configuring the edge server for the Master server 61
5.6.6 Configuring the edge server with a storage node 62
5.7 Performing additional tasks 65

1. About this guide

The *Support for Multiple Organizations* guide provides instructions for creating multiple organizations in a single inSync Private Cloud setup. This guide helps you understand the multi-org feature, and explains how to use the multi-org portal.

The *Support for Multiple Organizations* guide is primarily intended for Druva resellers.

Note: If you have any questions about the multi-org feature, contact the Druva Support team.

2. Overview of support for multiple organizations

2.1 Introduction

inSync Private Cloud supports the creation and management of multiple organizations (multi-org) via the Organization Portal. Organization Portal allows for easy creation and management of multiple customer accounts into a single installation of inSync Private Cloud. Organizations can now approach resellers to purchase accounts into such a ready Private Cloud setup. Using the Organization Portal, resellers can create multiple accounts, and assign these accounts to organizations. Because these accounts are pre-configured, consuming organizations can simply use the accounts to back up data. Using these accounts, organizations can enjoy an on-demand experience of Private Cloud. Multi-org support thus ensures easy management of on demand users of inSync Private Cloud. inSync resellers, who sell inSync as a service, typically deploy inSync Private Cloud within their premises. Using the Private Cloud setup, resellers create instances of Private Cloud and assign each such instance to customers. Customers can then use these instances, known as 'organizations', to access the Private Cloud setup. To each such organization, the Private Cloud setup appears dedicated. The multiplicity of instances ensures that the Private Cloud resources are distributed across organizations.

2.2 The multi-org benefits

Some of the top benefits of the multi-org feature are:

- **inSync Private Cloud as a service:** Using the multi-org feature, inSync resellers can offer inSync Private Cloud as a service. The user-friendly UI of the Organization Portal ensures that resellers can manage multiple organization accounts with ease. Besides, consuming organizations stand to benefit from the Private Cloud features such as scalability, enhances security, and superior backup performance.
- **No special configuration for multi-org portal:** The Organization Portal requires no special deployment or configuration; resellers only need to install and configure the Private Cloud setup. Accessing Organization Portal is as easy as modifying the inSync Master Management Console URL to include the path to the portal.
- **Intuitive UI:** The Organization Portal is simple to use, and intuitive in behavior. That is why resellers can get started with multi-org support smoothly while ensuring a flat learning curve.
- **Consolidated access to customer data:** The Organization Portal is a singular platform that allows resellers to create and manage multiple customer accounts. It offers a bird's eye view of the "status" of all organization accounts. Resellers do not require exhaustive efforts to get started; instead, they can create and manage multiple accounts with ease. Moreover, a single-platform approach to account management translates to huge savings in terms of administrative efforts and time.
- **No physical deployment necessary:** For organizations, the multi-org feature translates to freedom from investing in hardware otherwise required for an inSync Private Cloud setup. Rather, organizations enjoy Private Cloud as a service "on demand". Organizations simply log on to the Master server in a manner identical to logging on to an email service. This means that organizations can benefit from the Private Cloud experience without incurring the expenditure otherwise associated with the deployment of the setup.

2.3 How multi-org works

This diagram illustrates a simple implementation of multi-org support for two organizations.

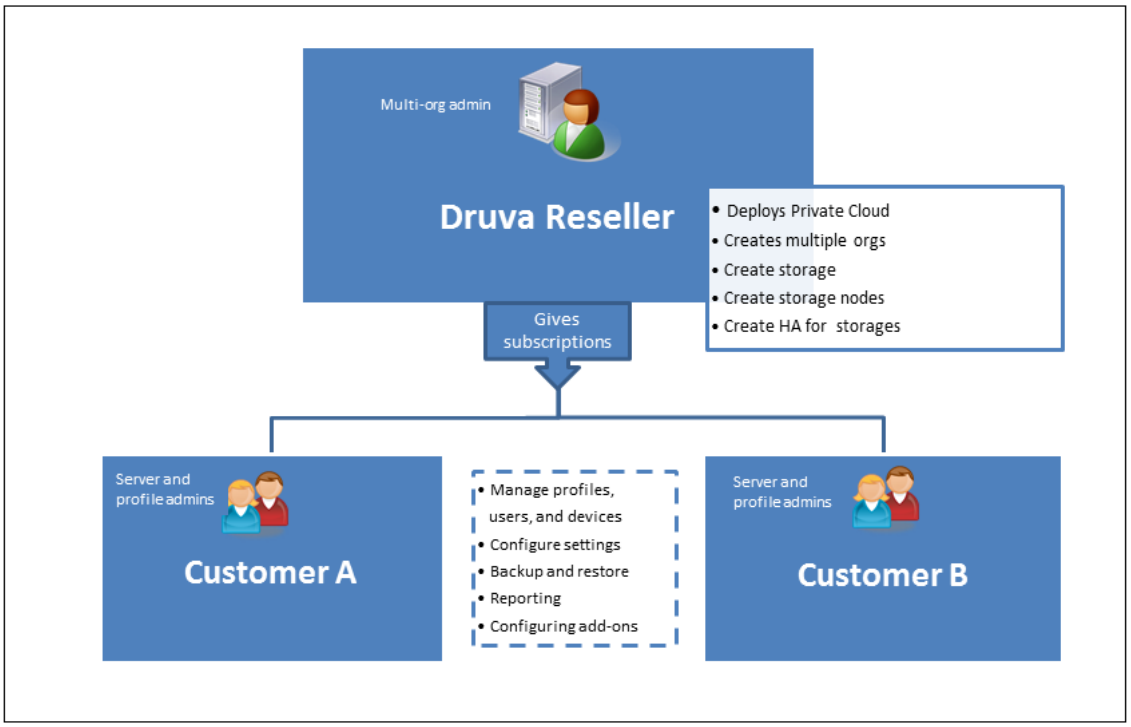

inSync resellers install Master server. The Master server installer also installs a local storage node on the same computer as the Master server. The behavior of this storage node is identical to that of the storage nodes that you create manually.

To configure the Private Cloud setup, the reseller uses the inSync Quick Configuration Wizard. This wizard ensures that one storage is created on the local storage node. Additionally, it also configures network and email settings for the Private Cloud setup.

Note: The initial configuration includes defining a first administrator for the Private Cloud setup. Because the reseller performs installation and configuration, the reseller becomes the first administrator of the Private Cloud setup.

To create multiple instances into the Private Cloud setup, the reseller logs on to the Organization Portal. The reseller creates multiple accounts for consuming customers from the Organization Portal. Each such account is an instance into the Private Cloud setup. The Organization Portal simplifies management of multiple such customer accounts. It allows a consolidated view of customer details such as name of the customer, type of Private Cloud license, add-on licenses and so on.
For this scenario, consider two customers, A and B, for whom the reseller creates two accounts.

Note: Adding an organization is identical to creating an account. The term "add" is used only to match with the Organization Portal UI.

For both A and B, the reseller performs initial configuration of accounts. While performing the configuration, the reseller considers factors such as size of the organization, number of users, and type of data for backup. These factors help the reseller make decisions such as creating the appropriate type of storage, assigning the correct size to the storage, and defining the high availability (HA) policy.

Thereafter, the inSync Private Cloud accounts are managed by administrators within the organization only. The reseller only prepares the inSync accounts for the customers A and B; the customers in turn, use these accounts for data backup and restore. Both a and B appoint administrators who perform tasks such as creating and managing profiles, creating and managing users and their devices, reporting, and configuring add-ons (Analytics, DLP, and Share).

Organizations A and B do not require a physical installation of inSync Private Cloud. Instead, they subscribe to an inSync account in a manner much similar to the creating accounts with email providers. inSync resellers can extend the multi-org support to as many organizations as required.

3. Planning for a Private Cloud setup

This topic contains a list of frequently asked questions around the Private Cloud deployment. These questions should help answer some of common queries related to the private cloud deployment.

Note: These FAQs are general guidelines that should help you plan your Private Cloud deployment. To fully understand your setup requirements, contact Druva Support.

Where do I install the Master server?

You can choose to install the Master server at a central location close to a majority of your storage nodes. As with all server-client applications connected via the internet, the Master server -- storage node communication depends upon latency issues. The latency might arise due to a geographical distance between the Master server and the storage node. To effectively back up data across your organization, the communication experience between the Master server and the storage nodes must be fair. That is why Druva recommends that your Master server should be centrally located as compared to storage nodes.

How many storage nodes do I create? Where do I create them?

You can create as many storage nodes as the number of data centers you have. The number of storage nodes depends on the number of customers you expect, and thus, the number of storage that you will create.

The number of storage nodes that you will need depends on the number of customers accounts (and hence the number of employees) that you create. One storage node can attend to approximately 2000 users. If you are planning to create a High Availability policy for your customers, you should plan to create at least 2 storage nodes for your setup.

How many storage do I create?

The Quick Configuration wizard creates a storage on the local storage node, which you can assign to customers. A storage node acts like a container for storage. However, data contained within a storage on a storage node cannot be accessed from other storage on the same node. That is why you can use

one storage node to create multiple instances of storage. Because data privacy across multiple storage on a storage node is guarded, you can create multiple storage for different customers on the same storage node.

You can consider creating new storage as the number of customer accounts grow. Each customer must be assigned to at least one storage. You might need to consider customer policies to decide how many storage you will require. For example, if a customer policy mandates that the senior management's data must be stored separately, create two storages for the customer - one for the senior management and one for the rest of the organization. You can create both the storage on the same storage node. Likewise, you can also create two storage on a storage node and assign these storage to different customers.

What should the size of each storage be?

It is important that you assign an optimum size to your storage. Because data resides in storage, the response time for backup or restore requests depends on the capacity of storage to handle such requests. If the size of your storage is not adequate, backup or restore performance might be affected. Alternatively, create large-sized storage for smaller volumes of data will only result in excessive investment in resources. The size of the storage depends on the data requirements of organizations for which you created accounts.

The size of each storage depends on factors such as:

- The number of users backing up data to the storage
- The average data each user will back up
- The daily incremental change in data for each user
- The type of data that is backed up
- The retention policy of the organization

Note: To estimate storage needs for each customer, use the ROI calculator available at http://www.druva.com/insync/roi-calculator/.

I want to create a High availability policy for my customers. What does this mean for storage nodes and storage?

The High Availability feature ensures that users enjoy a seamless backup or restore experience even if a storage is not available. To implement the High Availability feature, you must create a High Availability policy. To do this, you must create one secondary storage for each primary storage that you want to make "high available". You can configure the secondary storage for a scheduled backup. You can also ensure that data contained on the primary storage is seeded to the secondary storage periodically.

When you enable High Availability of a primary storage, you must create a secondary storage that is similar in size but residing on a different storage node as compared to the primary storage. This means that when you create a High Availability policy, you need at least two storage nodes and two storages for each policy.

Note: Factors such as actual number of primary and secondary storage, size of primary and secondary storage, and the seeding policy depend on individual customer requirements.

4. What you will need to get started

Before you use the Organization Portal to create and manage organizations, you must configure it. Whether you are performing a fresh deployment or upgrade, follow the chronology of this section to get your setup ready.

4.1 Installing the Master server

To access Organization Portal, you must install the Master server. For Windows specific installation instructions, see Installing the Master server on Windows. For Linux specific instructions, see Installing the Master server on Linux.

Note: The Master server installer also installs a local storage node. The behavior of this storage node is identical to the behavior of storage nodes that you create manually. However, you can only modify the name and the IP addresses of this storage node. To know how to do this, see Modifying a storage node.

4.2 Using the Quick Configuration Wizard

After you install the Master server, the Quick Configuration wizard appears. This wizard performs the following configuration:

- The URL and ports used by the Master server for backup and restore
- The email setup for the Master server
- The first file store on the local storage node

Note: The wizard creates a file store only. You can create an object store only after the initial configuration completes.

- The settings for the 'Default' profile.
- The resource consumption, access policies, notification policies, and retention policies for the 'Default' profile

- The first end-user

Note: The Quick Configuration Wizard ensures that a Private Cloud setup is ready for use. However, features such as the default profile and end-users cannot be configured from the Organization Portal. Although you have no control over such features, Druva recommends that you let the quick configuration complete. Thereafter, you can change network settings and email setup from the Organization Portal. You can also choose to create additional storage nodes, storage, storage pools, or HA policies.

For more information, see Using the Quick Configuration wizard.

4.3 Logging on to Organization Portal

This topic contains instructions for logging on to Organization Portal.

Before you begin, make sure that the following information is handy:

- The IP address for accessing Organization Portal

  Note: Use the IP address that you use for accessing inSync Master Management Console.

- Email ID and password of the first inSync administrator account

To log on to the Organization Portal

1. In a Web browser, enter the following URL:
   `https://<inSync Server URL>/orgportal`
   For example, `https://192.168.33.186/orgportal/`

   Note: The inSync Private Cloud setup supports HTTPS only. The Master server uses a self-signed certificate to authenticate an access to Organization Portal. The certificate ensures communication via HTTPS, but because it is self-signed, an untrusted certificate alert is displayed while accessing Organization Portal. Ignore this message and proceed with the login.

2. Enter the first administrator email ID and password.
3. Click **Login**.

4.4 Modifying your Private Cloud setup

The Quick configuration wizard ensures that your setup is ready for use. You can create multiple accounts, each an instance of the Private Clouds setup. You can however, choose to modify:

- The email setup
- The network settings

4.4.1 Modifying email setup

When you create or modify new organizations, the Master server sends emails to the respective administrators via the Organization Portal. To be able to send such emails, you require an SMTP server and an email account. The Quick Configuration wizard performs an initial email configuration for your setup. However, you can choose to modify these settings.

Before you begin, make sure that:

- You created an account that can be assigned to the Organization Portal.
- You have the SMTP server details.

To modify email setup

1. From the right-navigation pane, select **Settings** > **Email**.
2. Click **Edit**.
3. On the Email Setup window, modify:
   - **SMTP server**: The SMTP server
   - **SMTP port**: The port used by the SMTP server. The default value is 25.
   - **SMTP username**: The email ID that you want to assign to the Master server
   - **SMTP password**: The password for this email ID
   - **Use secure connection**: Select to enable TLS/SSL
4. (*Optional*) Click **Send Test Email** to test your settings.
5. Click **Save**.

4.4.2 Modifying the network settings

This section contains instructions for changing access ports and URLs, loading your SSL certificate, and managing IP addresses assigned to the Master server.

Changing access ports and URLs

The Master server uses these ports for data backup and restore:

- **Backup/sync port**: The port via which the Master server accepts client data. The default value is 6061.
- **Admin UI port**: The port used for accessing inSync Web. The default value is 443.
- **User web access URL**: The URL used for accessing inSync Web.
- **User activation URL**: The URL used for activating inSync clients.

To change the access ports and URLs

1. From the right-navigation pane, select **Settings** > **Network**.
2. Under Network Settings, click **Edit**. The **Network Settings** window appears.
3. Modify the ports and URLs.
4. Click **Save**.

Loading your own SSL certificate

inSync Private Cloud mandates communication via secure HTTP. If you do not load an SSL certificate, the Master server uses a self-signed certificate to authenticate access to Organization Portal. The certificate ensures communication via HTTPS, but because it is self-signed, an untrusted certificate alert is displayed during access. To resolve the untrusted certificate issue, upload an SSL certificate on the Master server.

Note: Druva recommends that in a production environment, you load an SSL certificate.

Before you begin, obtain an SSL certificate from a Certificate Authority and export it without the private key to a folder on your computer. Use a `.pem` certificate.

To load the SSL certificate

1. From the right-navigation pane, select **Settings** > **Network**.
2. Under Network Settings, click **Edit**. The **Network Settings** window appears.
3. Click the browse icon and select the certificate that you want to upload.
4. Click **Save**.

Note: You must now load the same certificate for all storage nodes. To know how to do this, see

Managing the IP address of Master server

You can configure a list of IP addresses or fully qualified domain names (FQDN) that for inSync clients to communicate with the Master server. You can also arrange these IP addresses or FQDNs in an order of your preference.

To add an IP address or FQDN

1. From the right-navigation pane, select **Settings** > **Network**.
2. Under Server IP / FQDN, click **Add New Server IP / FQDN**. The **Add Server IP / FQDN** window appears.
3. Enter the IP address / FQDN and the TCP port.
4. Set **Network type** to LAN, WAN, or AUTO.

   Note: If you set the Network type to AUTO, connections with time-to-live (TTL) 20 ms or higher are identified as WAN.

5. Click **Ok**.

Note: You can use the move up button to arrange the IP addresses in a desired sequence.

To change an IP Address or FQDN

1. From the right-navigation pane, select **Settings** > **Network**.
2. Under Server IP / FQDN, click the IP address you want to change and click **Edit**. The Edit Server IP / FQDN window appears.
3. Modify the **IP / FQDN** address, **TCP port**, or **Network type**.
4. Click **Ok**.

To delete an IP address

1. From the right navigation pane, select **Settings** > **Network**.
2. Under Server IP Address, select the IP address that you want to delete and click **Delete**.

5. Working with the Organization Portal

Follow the chronology of this section to work with the organization portal.

- Adding and managing organizations
- Creating additional storage nodes
- Creating additional storage
- Creating storage pools
- Creating HA policies
- Setting up inSync Edge server
- Performing additional tasks

5.1 Adding and managing organizations

When you add an organization, you create administrators and enable licenses for that organization. Thereafter, you should create storage and if required, HA policies for that organization.

Note: When you "add" an organization, you actually create an account for the participating organization.

5.1.1 Adding an organization

Follow the steps in this section to add an organization.

Before you begin, make sure that:

- You configured the Organization Portal.
- You created storage nodes.
- You gathered organization-specific information such as administrator names and licensing details.

To add an organization

1. From the right-navigation pane, click **Organizations**.
2. Click **Add New Organization**. The Add New Organization window appears.
3. On the first step of the wizard, specify:

   ***** Fields marked with an asterisk are mandatory.

   - ***Organization name:** The name of the organization
   - ***Contact name:** The name of the contact in the organization
   - ***Contact email ID:** The email ID of the contact in the organization
   - **Preferred time zone:** The time zone for the organization
   - **Phone:** The contact number of the organization
   - ***Administrator email IDs:** Email addresses of the administrators who will access the Master Management Console.

   Note: After you define a time zone, the timestamps on the inSync Master Management Console are displayed following the time zone.

4. On the second step of the wizards, specify:
   - **Organization type:** The type of license assigned to the organization
   - **Total users:** The number of users who can use inSync for data backup and share
   - **Total guests:** The number of guest users with whom inSync users can share data
   - **Total storage:** The size of the storage assigned to the organization
   - **License expires on:** The expiration date of the Private Cloud license
   - **Enable DLP, Enable Analytics, Enable share:** Indicates that add-ons are enabled

   Note: For each add-on that you select, specify an expiration date in the calendar box.

5. Click **Finish**.

5.1.2 Modifying an organization

Follow the steps in this section to modify organization details.

To modify an organization

1. From the right-navigation pane, click **Organizations**.
2. From the list of organizations, click the organization that you want to modify.
3. Under Summary, click **Edit** to modify **Organization name**, **Contact name**, **Contact email ID**, and **Phone**.
4. Save your changes.
5. Under Licensing Details, click **Edit** to modify license details.
6. Save your changes.

5.1.3 Enabling or disabling organizations

Follow the steps in this section to enable or disable an organization.

To enable or disable an organization

1. From the right-navigation pane, click **Organizations**.
2. From the list of organizations, click the organization that you want to enable or disable.

> **Note**: If you disabled an organization previously, the **Enable** button is enabled. Likewise, when you select an enabled organization, the **Disable** button is enabled.

3. To enable the organization, click **Enable**.
4. To disable the organization, click **Disable**.

5.1.4 Deleting an organization

When you delete an organization, the instance of Private Cloud associated with the organization is deleted.

Before you begin, make sure that all users of the organization and then the storage associated with the organization are deleted.

To delete an organization

1. From the right-navigation pane, click **Organizations**.
2. From the list of organizations, click the organization that you want to delete.
3. Click **Delete**.

5.1.5 Exporting user data

You can export customer data as a CSV file. The CSV will typically contain organization details that you entered while adding or modifying an organization.

To export user data

1. From the right-navigation pane, click **Organizations**.
2. Click **Export Data as CSV**.

Customer details are downloaded in a CSV file at a location that you specify.

5.2 Installing additional storage nodes

This section contains instructions for installing additional storage nodes.

5.2.1 Prerequisites for installing a storage node

This table contains a list of prerequisites to ensure that an installation of a storage node completes successfully.

Hardware prerequisites	
Prerequisite	**Minimum requirement**
CPU	64-bit, quad processor
RAM	8 GB for each TB of storage
Disk space	125 MB for installing the storage node
Software prerequisites	
Operating system	■ Windows 2008 R2 Service Pack 1 (64-bit) ■ Windows 2012 Server (64-bit) ■ Ubuntu 12.04 (64-bit) ■ Red Hat Enterprise Linux (RHEL) 6.3 (64-bit)

5.2.2 Installing a storage node

Follow the steps in this section to install a storage node on a Windows or Linux computer.

Before you begin, make sure that:

- You have administrator or root privileges on the computer on which you plan to install the storage node.
- You downloaded the inSync storage node installer.

To install a storage node on Windows

1. Double-click the storage node installer and click **Next**.
2. Accept the End-User License Agreement (EULA) and click **Next**.
3. In the **Choose Destination Folder** field, type or select the full path storage node home directory.
4. (*Optional*) Select shortcuts to access the storage node.
5. Click **Install** and then **Finish**.

To install a storage node on Linux

- For Ubuntu users

  From the directory that contains the .deb package, run the following command:

  ```
  sudo dpkg -i <package_name>
  ```

- For RHEL users

  From the directory that contains the .rpm package, run the following command:

  ```
  rpm -ivh <package_name>
  ```

Note: In these commands, <package_name> represents the file name of the installer.

5.2.3 Registering the storage node

For the Master server to identify and interact with the storage nodes, you must "create" the storage nodes from the Organization Portal.

Note: The actual process of creating a storage node involves Installing the storage nodes. Although called "creating", creating a storage node only involves establishing a connection between the storage node and the Master server.

Before you begin, make sure that:

- The following ports on the storage node remain free:
  - Backup and sync port: Port for sending backup data to the storage node. The default value 6061.
  - User web access port: Port for accessing inSync Web. The default port 443.

Note: You can enable communication via inSync Edge Server (edge server) for individual storage nodes. However, you cannot enable the edge server to validate client-server communication for individual organizations. Administrators of organizations must do so themselves. For a full set of documentation on edge servers, see Setting up edge servers.

To register the storage node

1. Generate the registration key.
   - (*Windows*): On the storage node server, select **Start** > **All Programs** > **Druva inSync Storage Node**.
     Click **Generate registration key**.
   - (*Linux*): Run the following command.
     ```
     sudo insync-storagenode-config.sh -k
     ```
2. Copy the registration key.
3. On the Organization Portal, from the right-navigation pane, select **Storage Nodes**.
4. Click **Create New Storage Node**.
5. On the first step of the wizard, enter:
   - **Storage node name**: The name of the storage node
   - **Primary IP (or) FQDN**: The IP address or FQDN of the storage node

- **Backup & sync port**: The port that the storage node uses to backup and synchronize data
- **User web access port**: The port used to access inSync Web
- **Storage node registration key**: The registration key of the storage node

6. (*Optional*) If you registered the edge server with the storage node, register the edge server with the Organization Portal. On the second step of the wizard, enter:
   - **Enable Edge Server**: Indicates that edge server association is enabled
   - **IP Address (or) FQDN**: The IP address or the fully qualified domain name (FQDN) of the edge server
   - **Edge Server Port**: The port used to connect with the edge server. The default port is 6061.

7. Click **Create Storage Node**.

5.2.4 Loading an SSL certificate on a storage node

If you loaded an SSL certificate on your Master server, you must load the same certificate across all storage nodes.

To load an SSL certificate

1. On the menu bar, click **Manage** > **Storage Nodes** > **<storage node name>**.
2. Under General Information, click **Load Certificate & Key**.
3. Select the certificate that you want to load.

5.2.5 Modifying a storage node

Follow the steps in this section to modify a storage node.

To change the storage node name

1. On the menu bar, click **Manage** > **Storage Nodes** > **<storage node name>**.
2. Under General Information, click **Edit**.
3. Enter a new name.
4. Click **Save**.

To modify ports

Note: Make sure you open new ports on the storage node.

1. On the menu bar, click **Manage** > **Storage Nodes** > <**storage node name**>.
2. Click the storage node for which you need to change ports.
3. Under General Information, click **Edit**.
4. Enter the new ports.
5. Click **Save**.

To add an IP address or FQDN

1. On the menu bar, click **Manage** > **Storage Nodes** > **<storage node name>.**
2. Under IP Address or FQDN Configuration, click **Edit**.
3. Click **+Add IP Address or FQDN**.
4. Enter the new IP address and select the network type.

Note: You can set the network type as LAN, WAN, or AUTO. If the Server IP address is set to AUTO, connections with time-to-live (TTL) 20 ms or higher are identified as WAN.

5. Click **Save**.

To select a primary IP address

Note: The primary IP address is the first address that the Master server uses to communicate with the storage node. The Master server uses subsequent IP addresses (if any) if communication via the primary IP address fails.

1. On the menu bar, click **Manage** > **Storage Nodes** > **<storage node name>**.
2. Under IP Address or FQDN Configuration, click **Edit**.
3. Select the IP address that you want to configure as the primary IP address.
4. Click **Save**.

To change an IP address or FQDN

1. On the menu bar, click **Manage** > **Storage Nodes** > **<storage node name>**.
2. Under IP Address or FQDN Configuration, click **Edit**.
3. Change the IP address and set the network type.
4. Click **Save**.

To delete a non-primary IP address or FQDN

Note: You cannot delete the primary IP address.

1. On the menu bar, click **Manage** > **Storage Nodes** > **<storage node name>**.
2. Under IP Address or FQDN Configuration, click **Edit**.
3. Click the delete icon next to the IP address you want to delete.
4. Click **Save**.

5.3 Creating and managing additional storage

When you perform an initial configuration of your Master server,

5.3.1 Understanding types of storage

inSync supports two types of storage:

- Object store implemented by using OpenStack Swift in conjunction with OpenStack Keystone

Note: To fully understand the Swift and Keystone architecture, see the Swift documentation available at http://docs.openstack.org/developer/swift/. The Keystone documentation is available at http://docs.openstack.org/developer/keystone/.

- File store

Object store

An object store is a type of storage that is designed to treat data as object. Each object contains data and metadata, and is associated with a unique identifier. Numerous such objects are arranged in a flat hierarchy and are typically accessed using some API.

inSync supports the creation of object stores, provided that you implement it using OpenStack Swift (Swift). The Swift object store arranges data within "containers" designated for such a purpose. To know more about OpenStack Swift, see https://wiki.openstack.org/wiki/Swift.

When you create an inSync object store, inSync creates a container within Swift. This container holds backup data. inSync creates such containers for each object store that you create.Object stores have the following limitations:

- You cannot create additional data volumes on an object store.
- If you associate an object store with users, you cannot change the storage for these users.
- Object stores cannot belong to a storage pool.

File store

As the name suggests, a File store is a sequence of bytes of a length that is structured like a block. A File store treats data as blocks within sectors and tracks of a storage media. You can think of each such block of data as an individual hard drive, and control it just the way you would control a hard drive. Managing storage of the type File store is simple; you can attach such storage to an external

server. The external control of this storage control matches with your inSync Private Cloud architecture thus ensuring seamless storage management.

A File store is especially suited for environments where storage requirements grow over time. Because a Private Cloud setup witnesses high data volumes, you can add multiple data volumes to your File store whenever required. A File store type of storage thus comes with the benefit of scalability.

A File store can be extended to contain multiple to which backup data is saved. Data volumes are discrete units of storage on a physical disk that supports data. Because a File store includes multiple data volumes, it becomes a "cluster" of storage media to which data is saved.

Note: You cannot control how data is saved to data volumes. After a data volume is filled to its capacity, the subsequent data is saved to the next data volume of that storage.

How file store and object store differ

This table explains how file store and object store are different.

Property	File store	Object store
Deduplicated data stored in...	File system	Object store **Note:** inSync supports object store implemented with OpenStack Swift in conjunction with OpenStack Keystone only.
Increase storage capacity by...	Adding up to 10 data volumes	Increasing the size of the storage. **Note:** Before you increase the size of your object store, make sure that your Swift database contains enough space.
Maximum size	64 TB	64 TB **Note:** A single installation of Swift can contain multiple inSync storages.
Support for storage pools	Yes	No

5.3.2 Creating a file store

This section contains instructions for creating a file store on Windows and Linux computers.

Note: By default, storage is optimized for performance. You cannot optimize storage for "Disk space savings". However, inSync 5.4 supports the storage that you created and optimized for "Disk space savings" in earlier deployments.

Best practices for creating File store folders

If you expect large volumes of data for backup or restore, consider creating storage of the type File store. Because you can add multiple data volumes to your File store, File stores are scalable in nature.

Each File store contains three folders:

- Data folder - It contains the backup data.
- Database folder - It contains the inSync database, which contains the metadata of the backup data.
- Database log folders - It contains log files for all inSync database activities.

Follow these best practices while creating Data, Database, and Database log:

- Create the Database and Database Log folders on a local drive. Do not create these folders on a networked or shared drive.
- Create the Data folder on a local drive, Storage Area Network (SAN), or a Network-attached Storage (NAS).
- Create the Database folder on a separate disk. If you expect many users to back up their data on the storage, create the Database folder on a Solid-state Drive (SSD).
- Create the Data folder and the Database log folder on separate disks.

Before you begin, make sure that:

- You created the organizations that you want to associate with the storage.

Note: You can attach a storage with one organization only. However, you can attach multiple storage with an organization.

- The storage node on which you want to create the storage is up and running.
- You have administrator or root privileges on the storage node.
- Your File store folders (Data, Database, and Database log) are excluded from an antivirus scan.
- (*Linux*) You created folders that you want to use as Data, Database, and Database log folders on your storage node.
- (*Linux*) You assigned the user 'insyncuser" group and owner permissions on Data, Database, and Database log respectively.

  To assign group and owner permissions, use the chown command in the following format:

  ```
  $ sudo chown insyncserver:insyncserver <storagefolder>
  ```

To create a File store

1. On the menu bar, click **Manage** > **Storage**.
2. Click **Create New Storage**. The storage creation wizard appears.
3. Enter the relevant details on each step of the wizard. For information, see Fields in the storage creation wizard.
4. Click **Finish**.

Fields in the storage creation wizard

This section contains a description of the fields that appear when you create storage.

General Information

This table describes the fields that appear on step 1 of the wizard.

Field	Description
Storage name	The name of the storage.
Storage node	The storage node on which you want to create the storage.
Storage type	The type of storage (primary or secondary) that you are creating.
Attach to organization	The organization to which you want to attach the storage.

Data Storage Details

This table describes the fields that appear on step 2 of the wizard.

Field	Description
Where do you want to store data?	Select **File Store**
Data folder	The Data folder of the first data volume.
Size	The size of data that can be stored in the storage.
Storage consumption alert threshold	The storage consumption threshold value, which when crossed, will trigger a low disk space alert. For example, if the Max. disk space for the storage is 10 TB and storage consumption alert

	threshold is 80%, an alert is sent to the administrators when storage consumption exceeds 8 TB.

Performance

This table describes the fields that appear on step 3 of the wizard.

Field	Description
Database folder	The Database folder or the SSD path of the storage.
Database log folder	The Database log folder.
Max. parallel connections	The maximum number of parallel connections to the storage. **Note:** You can set this field to a maximum value of 500.

Compaction Schedule

This table describes the fields that appear on step 4 of the wizard.

Field	Description
Compact daily at	The start time and duration for daily storage compaction.
Compact weekly on	The scheduled day for weekly compaction of the storage.
Compact weekly at	The start time and duration for weekly storage compaction.

5.3.3 Creating an object store

This section contains instructions for creating an object store on Windows and Linux.

Best practices for creating object store folders

If you implemented OpenStack Swift (Swift) in conjunction with Keystone in your organization, you can use the Swift object store to create object store type of storage. Object stores are especially compatible with a Private Cloud setup, where connectivity and latency might fluctuate.

Each object store type of storage that you create will contain:

- Database folder - It contains the inSync database, which contains metadata of the backup data.
- Database log folders - It contains log files for all inSync database activities.

Follow these best practices while creating Database and Database log:

- Create the Database and Database Log folders on a local drive. Do not create these folders on a networked or shared drive.
- Create the Database folder on a separate disk. If you expect many users to back up their data on the storage, create the Database folder on a Solid-state Drive (SSD).

Before you begin

- If you are implementing Swift with Key stone especially for your Private Cloud setup, make sure that you configured your Swift setup.
- Make sure that you have Swift details such as IP address, port, access key, and pass key.
- Your object store folders (Database and Database log) are excluded from an antivirus scan.
- (*Linux*) Make sure that you created folders that you want to use as Database and Database log folders on your storage node.
- (*Linux*) You assigned the user 'insyncuser" group and owner permissions on Database and Database log respectively.

To assign group and owner permissions, use the chown command in the following format:

```
$ sudo chown insyncserver:insyncserver <storagefolder>
```

To create an object store

1. From the menu bar, click **Manage** > **Storage Lists**.
2. Click **Create New Storage**.
3. Enter the relevant details. For more information, see Fields on the storage creation wizard.
4. Click **Finish**.

Fields on the storage creation wizard

This section contains a description of the fields that appear while creating a storage.

General Information

This table describes the fields that appear on step 1 of the wizard.

Field	Description
Storage name	The name of the storage.
Storage node	The storage node on which you want to create the storage.
Storage type	The type of storage (primary or secondary) that you are creating.

Attach to organization	The organization to which you want to attach the storage.

Data Storage Details

This table describes the fields that appear on step 2 of the wizard.

Field	Description
Where do you want to store data?	Select **Object Store.**
Object store IP/FQDN	The IP address or the fully qualified domain name (FQDN) of the object store.
Port	The port for connecting to the object store. The default value is 8080.
Access key	The access key for Keystone.
Secret key	The pass key for Keystone.
Size	The size of the object store.
Storage consumption alert threshold	The storage consumption threshold value, which when crossed, will trigger a low disk space alert. For example, if the Max. disk space for the storage is 10 TB and storage consumption alert threshold is 80%, an alert will be sent to the administrators when storage consumption exceeds 8 TB.

Performance

This table describes the fields that appear on step 3 of the wizard.

Field	Description
Database folder	The Database folder or the SSD path of the storage.
Database log folder	The Database log folder.
Max. parallel connections	The maximum number of parallel connections to the storage. **Note**: You can set this field to a maximum value of 500.

Compaction Schedule

This table describes the fields that appear on step 4 of the wizard.

Field	Description
Compact daily at	The start time and duration for daily storage compaction.
Compact weekly on	The scheduled day for weekly compaction of the storage.
Compact weekly at	The start time and duration for weekly storage compaction.

5.3.4 Modifying storage details

Follow the steps in this section to modify storage details.

To modify a storage

1. From the right-navigation pane, click **Storage > <storage name>**.
2. Click **Summary**.
3. Under **Summary**, click **Edit** to modify **Storage name,** and save your changes.
4. Under **Data Storage Details**, click **Edit** to modify the data volume **Size** and **Storage consumption alert threshold,** and save your changes.

Note: If you are modifying an object store, **Size** indicates the size of storage.

5. Click **Performance & Compaction**.
6. Under **Storage Compaction Schedule**, click **Edit** to modify compaction details, and save your changes.
7. Under **Performance**, click **Edit** to modify **Max. parallel connections,** and save your changes.

5.3.5 Managing data volumes on a storage

You can associate additional data volumes with File store type storage.

Before you begin, make sure that:

- (*Linux*) You created the Data folder on the storage media that will be used as the data volume.
- Set the user "insyncserver" as the owner, and assign the group "insynserver" to this user. Use the chown command in the following format:

```
$ sudo chown insyncserver:insyncserver <Data folder>
```

To add a data volume

1. From the right-navigation pane, click **Storage > <name of File store>**.
2. Under **Summary** > **Data Volumes**, click **Add Data Volume**
3. In the **Data folder** box, select the Data folder for the data volume.
4. In the **Size** box, enter size of the data volume (expressed in GB or TB).
5. Click **Add Data Volume**.

To modify data volume size

1. From the right-navigation pane, click **Storage** > <**name of File store**>.
2. Under **Summary** > **Data Volumes**, click **Edit**.
3. In the **Size** box, set the size of the data volume (GB or TB).
4. Click **Save**.

To delete a data volume

You cannot delete individual data volumes. To delete a data volume, you must delete the storage. For instructions, see Deleting a Storage.

5.3.6 Deleting a storage

Follow the steps in this section to delete a storage.

Before you begin, make sure that you deleted all users assigned to the storage.

To delete a storage

1. From the right-navigation pane, click **Storage**.
2. From the list of storage, click the storage that you want to delete and click **Delete**.

5.3.7 Compacting a storage

Compaction is a process of deleting expired data to make room for new data. With each backup, the previous backup data becomes obsolete. This obsolete data is retained for a period that your retention policy dictates. For example, if your organization's policy mandates a retention of 90 days, the "obsolete data" is retained for 90 days. At the end of 90 days, this data is deleted from the storage, thus freeing up storage space.

The compaction process is resource-intensive and time-consuming, and thus, should run during off-peak hours. Druva recommends that along with the daily compaction, you run a weekly compaction for a longer duration.

Note: Compaction and user migration do not run simultaneously. If a user migration is in progress, compaction is temporarily halted. Similarly, if compaction is in progress, automatic user migration is temporarily halted.

To compact a storage

1. From the right-navigation pane, click **Storage**.
2. From the list of storage, click the storage to which you want to compact and click **Compact Now**.

To modify the compaction schedule of a storage

1. From the right-navigation pane, click **Storage** > <**storage name**>.
2. Under **Performance & Compaction** > **Storage Compaction Schedule**, click **Edit**.
3. On the Edit Storage Compaction window, enter:
   - **Compact daily at**: The start time and in the accompanying textbox, a new duration.
   - **Compact weekly on**: A day for weekly compaction.
   - **Compact weekly from**: The start time and in the accompanying textbox, a new duration.
4. Click **Save**.

5.4 Creating storage pools

This section contains instructions for creating storage pools.

5.4.1 Understanding storage pools

A storage pool is a collection of file stores that function as several units forming a whole. Each such unit backs up data. In large deployments where large volumes of data flow within an organization, the maximum stretchable limit of storage size (64 TB) might be insufficient. By pooling large-sized storage, inSync Private Cloud functions as a "big storage" for large deployments. Besides, the capability to dynamically add storage to a pool ensures scalability and support for organizations anticipating increasing data needs.

Some benefits of implementing a storage pool are:

- By distributing users across storage in proportion to the free space available on them, a pool acts like a load-balancer
- By automatically migrating users and their data from an occupied storage to a relatively empty one, a pool acts like a load-balancer.
- By allowing support for adding multiple storage at any time, a storage pool acts like a scalable resource.

Load-balancing by distributing users

When you create users, you can assign them to storage pools or storages within the pool. If you assign users to a pool, user data is distributed across member storages. The size of storage reserved for each user depends on the free space available across a storage. A storage with more free space will be assigned to store more user data. However, if you assign users to a storage within the pool, user data is backed up to this storage only.

Note: The Low Database Storage Available alert is generated for individual storages. This means that an alert is generated for every storage within a pool exceeding its consumption threshold.

Load-balancing by migrating users

When a storage in a pool is filled to 80% of its capacity while other storages are relatively free (occupied by up to 70% only), the storage pool becomes "unbalanced". New data requests are not assigned to the storage having less than 30% free space; instead, users assigned to it are automatically migrated to the freer storages. The migration of users ensures freeing up of space on the source storage while ensuring fair distribution of load across other relatively free storages. To understand how migration works, see How user migration in a storage pool works.

A storage pool is a collection of storage to which data is backed up. A storage pool functions as a consolidated view of several storage that act as backup for each other. Storage pools ensure adequacy of storage space for large volumes of data. Because a pool contains several storage, data can "move between" storage in scenarios where some storage are overwhelmed with data. This means that backup and restore activities are evenly distributed, thus ensuring that no one storage is overloaded, while the others remain free. Storage pools come with the benefit of storage scalability and load balancing.
Some benefits of implementing a storage pool are:

- Storage-level load balancing for heavy data flow.
- Automatic storage failover (if a storage becomes unhealthy, the system seamlessly keeps running).
- Ease of administration; you can easily manage multiple pools from the inSync Master management Console.

5.4.2 How user migration works

You can add multiple storage to a storage pool. Storage within this pool respond to backup and restore. However, when a storage in a pool is filled to 80% of its capacity while other storage are relatively free (occupied by up to 70% only), the storage pool becomes "unbalanced". New data requests are not assigned to the storage having less than 30% free space; instead, users assigned to it are automatically migrated to the freer storage. User migration ensures fair distribution of load between member storage.

Note: User migration is an automatic process.

This diagram explains how data migration works within a pool.

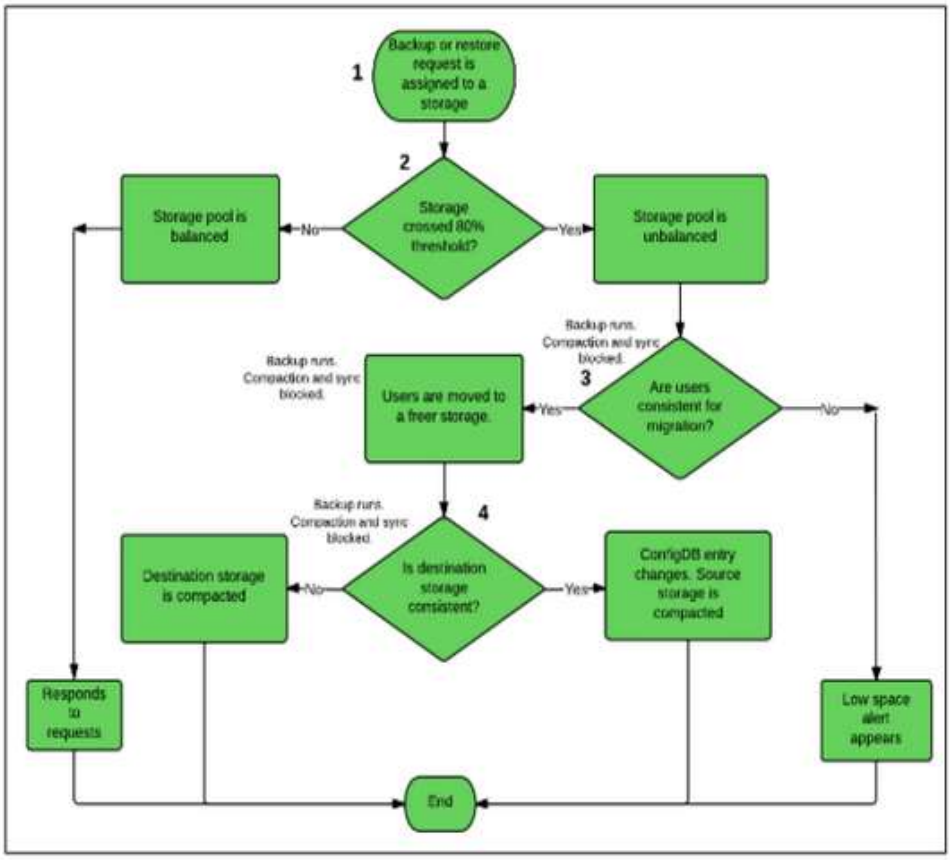

Step 1: The Master server assigns a backup or a restore request to a storage.

Step 2: If storage in a pool is filled to 80% of its capacity while other storages are relatively free (occupied by up to 70% only), the storage pool becomes "unbalanced". Backup and restores requests to this storage stop, and it is marked for migration.

Step 3:

1. The Master server initiates a scan on the storage to identify users that can be migrated to other storages within the pool. If the scan identifies consistent users, these users are moved to along with their snapshots a free storage, thus freeing up space on the source storage. The storage pool returns to a balanced state.
2. If the Master server cannot identify consistent data, it generates a Low Free Space Available alert.

Note: Backups continue to progress during this time. However, compaction and synchronization are temporarily halted.

Step 4:

1. The Master server initiates a scan on the destination storage to determine if the newly migrated users are consistent. If users are consistent, the Master server changes the ConfigDb entry to reflect changed storage. At this stage, the source storage is compacted.
2. If the data is inconsistent, a compaction runs on the destination storage to remove invalid data entries.

Note:

- User migration fails if consistency checks described in step 3 and step 4 do not succeed.
- At the end of successful migration, user data is backed up to the new storage.
- A status of migration activities is saved to inSyncMigration.log. This file is created on storage nodes.

5.4.3 Creating a storage pool

You can create storage pools at any time. Each such pool can contain multiple storage.

Note: You cannot associate object stores with a pool.

Read this first

- If you expect large volumes of data, create additional storage nodes before creating a storage pool.
- You must create storage before creating a pool. If you select to create a pool first, the storage pool creation wizard redirects you to the storage creation wizard.
- A storage can belong to only one pool.
- Only primary storage can belong to a storage pool. However, you can associate these primary storage with secondary storage thus ensuring high availability of the primary storage.
- You can continue to create data volumes on File stores that belong to a pool.
- Flexible user-storage mapping is not supported between storage belonging to a pool. Likewise, you cannot change user storage to a storage that belongs to a pool.

To create a storage pool

1. From the menu bar, select **Manage** > **Storage Pools**.
2. Click **Create New Storage Pool**.
3. On the Create New Storage Pool window, enter a **Storage pool name**.
4. In the **Storage to be added** list, select the storage that you want to associate with the pool.

   Note: The available list of storage only contains primary storage that do not belong to another pool.

5. Click **Create Storage Pool**.

   Note: An entry appears in the **Admin Audit Trail** for each storage pool that is created.

5.4.4 Modifying storage pool

Follow these steps to modify a storage pool.

To modify storage pool name

1. From the menu bar, click **Manage** > **Storage Pools** > **<storage pool name>**.
2. Under General Storage Pool Information, click **Edit**.
3. In the Edit Storage Pool Summary window, modify the Storage pool name.
4. Click **Save**.

To associate a file store with a pool

1. From the menu bar, click **Manage** > **Storage Pools** > **<storage pool name>**.
2. Under Assigned Storage, click **Attach Storage**.
3. In the Storage to be added list, select the storage that you want to associate with the pool.

Note: The available list of storage only contains primary file stores that do not belong to another pool.

4. Click **Attach**.

To disable data migration

1. On the Master server, locate the inSyncServer.cfg file.
   (*Windows*) C:\inSyncCloud\inSyncServer4
   (*Linux*) /etc/inSyncCloud/inSyncServer
2. Open the inSyncServer.cfg file in a text editor.
3. Locate the MIGRATION_ENABLED parameter and set it to False.
4. Save your changes.

Note: If a migration is in progress when you disable data migration, the on-going process does not stop. Instead, data migration stops thereafter. A status of data migration activities is saved to inSyncMigration.log. This file is created on storage nodes.

To delete a storage pool

Note: Before you begin, make sure that no backups to the pool are running.

1. On the menu bar, click **Manage** > **Storage Pools**.
2. Click the storage pool that you want to delete and click **Delete**.

Note: When you delete a storage pool, storage belonging to that pool are "released". Such storage continue to exist, but stop being a part of the deleted pool.

5.5 Creating an HA policy

Creating a high availability (HA) policy ensures that your primary File store is always available for data backup and restore. When you ensure high availability of primary storage, data contained on the primary storage is replicated on the secondary storage. In event of a failure of the primary storage, the secondary storage can handles backups and restores. Users do not witness any change in backup or restore experience. In event of primary storage failure, the secondary storage handles backup or restore requests. The primary storage is thus, "highly available".

Note: Druva recommends that you create an HA policy for File stores. If you plan to create an HA policy for object stores, contact Druva Support first.

Before you begin, make sure that:

- You created a secondary storage of the same size and configuration as the primary storage to which you want to associate the secondary storage.
- The secondary storage is on a different storage node than the primary storage.

To create a high availability policy

1. From the right-navigation pane, click **HA** and then **Create New**.
2. On the first step of the wizard, enter:
   - **Policy name**: Name of the HA policy
   - **Primary Storage**: Primary storage for the policy
   - **Secondary Storage**: Secondary storage for the policy
   - **Backup to Secondary every**: Frequency of backup from user devices to the secondary storage.
3. On the second step of the wizard, enter:
   - **Enable Seeding**: Enable seeding from primary storage to secondary storage.

   Note: The other fields appear only if you select Enable Seeding.

   - **Populate secondary storage daily at**: Duration for seeding.
   - **Additionally, populate weekly on**: Day on which weekly data is seeded.

   Note: Daily seeding is not be triggered on the day weekly seeding is scheduled. The seeding stops after the duration you specify even if all changes are not seeded.

   - **Populate secondary storage weekly at**: Weekly seeding schedule.
   - **Use dedicated IP/FQDN for destination**: Enable a dedicated IP or fully qualified domain name to seed the secondary storage.

   Note: If you do not select this check box, you can select an IP address from the list of IP addresses in the IP/FQDN for destination field.

   - **IP/FQDN for destination**: The IP address or FQDN for seeding the secondary storage.
   - **Limit bandwidth consumption to**: The bandwidth for seeding. Enter 0 for unlimited bandwidth. 0 KBps or 0 MBps.
4. Click **Finish**.

5.6 Setting up inSync Edge Server

This section contains instructions for setting up inSync Edge Server.

5.6.1 Understanding the edge server

inSync Edge Server (edge server) enables inSync clients located outside of your organization's firewalls to communicate with the Private Cloud setup. As the name suggests, it is a server that "sits between" the inSync network and the Internet, typically in a demilitarized zone (DMZ). By validating all internet-facing data flow, the edge server introduces an additional layer of security in the Internet -- inSync network communication. inSync clients using WAN can only communicate with the Private Cloud setup after the edge server validates. Likewise, if your storage nodes are outside the Master server premises, the edge server validates communication between the Master server and the storage nodes.

The edge server ensures that the Master server can respond to client requests smoothly; by validating client requests, it ensures that the Master server is not overwhelmed with unacceptable requests. It also acts as a gateway that filters verified requests via "unverified" networks, thus counteracting the vulnerability of the Private Cloud setup.

5.6.2 Understanding inSync Private Cloud architecture

This diagram depicts a simple deployment of inSync Private Cloud with edge server.

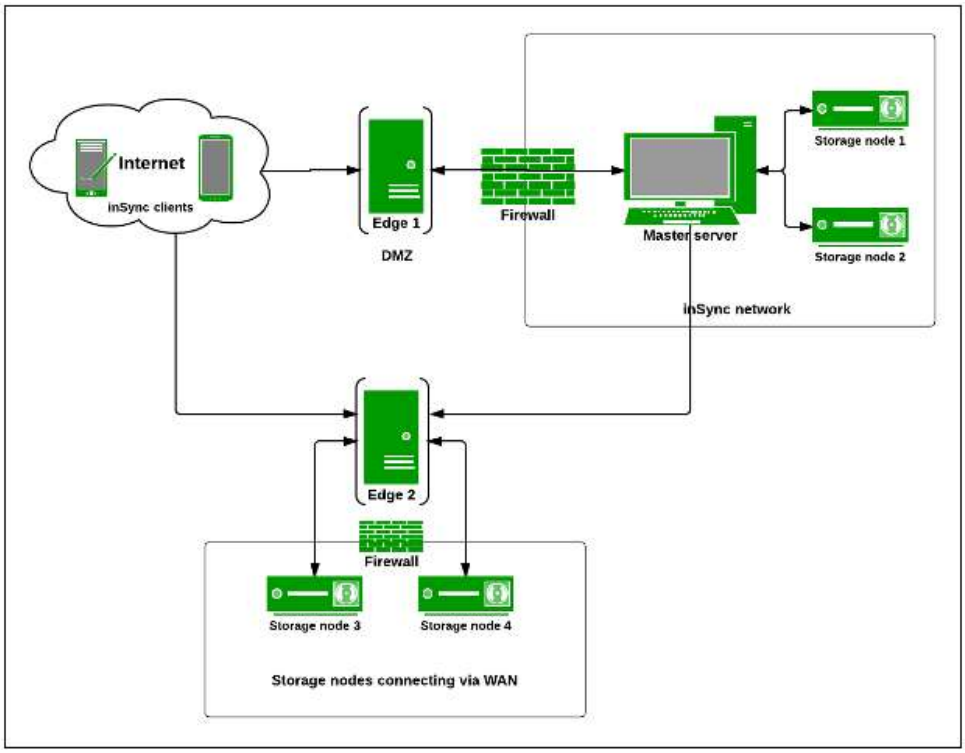

Name	Description
inSync clients	inSync clients, which are installed on endpoints, initiate backup or restore requests via LAN or WAN. inSync client requests via WAN are validated by the edge server, and then authenticated by the Master server. The Master server routes authenticated requests to storage nodes. **Note**: inSync clients, once validated by the edge server, "inform" the Master server about a backup or restore requirement. The clients do not route any data to the Master server. The Master server only acknowledges and assigns client requests to storage nodes. inSync clients send or receive data to or from storage nodes. The edge server only "validates" and the Master server only "assigns". The storage nodes actually back up or restore data.
Edge server	The edge server, which sits on the "edge" of the inSync network, performs a one-time validation of requests originating from the internet. The Master server or storage nodes perform a "handshake" with the edge server, thus authorizing it to validate client requests via WAN. The edge server, after validating such requests, creates a channel of communication between each client and the Master server. If the storage node is located outside LAN limits, a separate instance of edge server creates a channel between clients and storage node as well as the Master server and storage node. **Note**: The edge server configured with the Master server also authenticates client requests to the local storage node and the storage nodes within LAN limits. In this scenario, the Master server requests to storage node do not require authentication. However, if the storage nodes are deployed outside LAN limits, a

	separate installation of edge server validates all Master server and client requests. **Note**: You can configure the edge server to work with a combination of Master server and storage nodes, provided that they are located within LAN limits.
Master server	The Master server, which is centrally located, manages storage nodes as well as user information. The Master server validates, authenticates, and assigns backup or restore requests routed over LAN and assigns them to storage nodes. If client requests are routed via WAN, the edge server validates them, and creates a communication channel between clients and Master server. Thereafter, the Master server uses this channel to authenticate and assigns such requests to storage nodes. **Note**: You can configure your Master server to recognize only one edge server.
Storage nodes	A storage node is a server that performs backups or restores. A collection of such storage nodes is nothing but a private cloud. In a Private Cloud deployment, multiple storage nodes are located within and outside of the Master server perimeter. To ensure added security, the storage node performs a "handshake" with the edge server, thus permitting it to validate incoming requests. If a storage node is geographically co-located with the Master server, the same instance of edge server performs validation of client requests for the storage node and the Master server.

However, if a storage node communicates with the Master server via WAN, a separate instance of edge server validates both client and Master server requests to the storage node. Thereafter, the edge server creates a communication channel between each client and the storage node as well as the Master server and the storage node.

Note: You can configure your storage node to recognize only one edge server.

5.6.3 How inSync Private cloud works

This section explains the workflow of an inSync Private Cloud deployment with an edge server.

Step 1: inSync client sends a backup or restore request to Master server via WAN.

Step 2: This request is intercepted by the edge server. The edge server validates such a request, and if the validation is successful, creates a communication channel between the client and the Master server. The client and the server now use this channel for communication; the edge server does not intercept subsequent data flow.

Note: The Master server must first perform a "handshake" with the edge server. To know how this is done, see Configuring the edge server for the Master server.

Step 3: The Master server acknowledges and authenticates the client request. It assigns the request to an appropriate storage node.

Step 4:

1. If the storage node is located within the Master server perimeter, the edge server does not validate the Master server request to the storage node.
2. If the storage node is located outside of Master server network, the edge server configured with the storage node validates Master server request.

Step 5:

1. The edge server validates the client request for data backup or restore. If the validation is successful, the edge server creates a communication channel via which inSync client sends data to the storage node. The edge server does not intercept subsequent communication between the client and the storage node.

Note: The same installation of edge server performs validation for Master server and storage nodes located within LAN limits.

2. The edge server validates the client request for data backup or restore. If the validation is successful, the edge server creates a communication channel via which inSync client sends data to the storage node. The edge server does not intercept subsequent communication between the client and the storage node.

Note: For storage nodes located outside of Master server perimeter, a separate installation of edge server performs validation of client requests.

Step 6: inSync client sends data to the storage node. Alternatively, the storage node restores data to desired location.

5.6.4 Installing the edge server

Follow the platform specific steps to install the edge server.

Installing the edge server on Windows

Before you begin, make sure that your system adheres to the following minimum requirements.

Hardware	
Requirement	Specification
CPU	64-bit, dual core processor
RAM	8 GB
Network card	10/100 Mbps fast Ethernet adapter
Disk space	125 MB for installing edge server
Software	
Requirement	Specification
Operating system	■ Windows 2008 R2 Service Pack 1 (64-bit) ■ Windows 2012 Server (64-bit)
Web browser	■ Internet Explorer 8 or later ■ Firefox 19 or later ■ Safari 5.1 or later ■ Chrome 25 or later

To install the edge server

1. Download the edge server installer.
   a. In the right navigation pane, click **Download Edge Server**.

Note: To ensure that your installation completes, make sure that you have administrator privileges.

2. Double-click the edge server installer and click **Next**.
3. Accept the End-User License Agreement (EULA) and click **Next**.
4. In the **Choose Destination Folder** box, type or select the full path to the installation home directory.
5. (*Optional*) Select shortcuts for launching the edge server.
6. Click **Install** and then **Finish**.

Installing the edge server on Linux

Before you begin, make sure that your system adheres to the following minimum requirements.

Hardware	
Requirement	Specification
CPU	64-bit, dual core processor
RAM	8 GB
Network card	10/100 Mbps fast Ethernet adapter
Disk space	125 MB for installing edge server
Software	

Requirement	Specification
Operating system	■ Ubuntu 12.04 (64-bit) ■ Red Hat Enterprise Linux (RHEL) 6.3 (64-bit)
Web browser	■ Internet Explorer 8 or later ■ Firefox 19 or later ■ Safari 5.1 or later ■ Chrome 25 or later

To install the edge server

1. Download the edge server installer.
   a. In the right navigation pane, click **Download Edge Server**.

Note: To ensure that your installation completes, make sure that you have root privileges.

2. Install the edge server.
   (*Ubuntu*) From the directory that contains the .deb package, run the following command:
   `sudo dpkg -i <package_name>`

   (*RHEL*) From the directory that contains the .rpm package, run the following command:
   `rpm -ivh <package_name>`

Note: In these commands, <package_name> represents the filename (along with the extension) of the installer.

5.6.5 Configuring the edge server for the Master server

After you install the edge server, you must configure it to work with the Master server. When you configure the edge server, you establish a "connection" between the Master server and the edge server via which inSync client requests authenticated by the edge server are routed. Registering the edge server thus ensures that the "handshake" between the Master server and the edge server is complete.

Before you begin, make sure that:

- You installed the Master server.
- You can log on to the computer on which the edge server in installed.
- Port 6061 on the edge server remains free. inSync clients use this port to connect with the edge server.

Step 1: Generate the edge server key

The registration key is required for establishing a connection between the Master server and edge server. It acts as a unique identifier and permits the Master server to identify and interact with the edge server.

To generate an edge server key

1. Log on to the edge server.
2. Generate the registration key:

   Windows: From the **Start** menu, click **All Programs** > **Druva inSync Edge Server**.

   Linux: Run the following command.

   ```
   sudo insync-edgeserver-config.sh –k
   ```
3. Copy the registration key.

Note: Druva recommends that you generate the registration key only once. If you generate the registration key again, the old registration key becomes inactive. If you generate the key multiple times, register the edge server using the most-recent key.

Step 2: Register the edge server

Follow the steps in this section to register the edge server with the Master server.

To register the edge server

1. Log on to the inSync Master Management Console.
2. From the menu bar, select **Manage** > **Settings**> **Edge Servers**.
3. On the Edit Settings window, enter the following details:

   - Enable edge server for my master: Indicates that edge server association is enabled
   - IP Address/FQDN: The IP address or the fully qualified domain name (FQDN) of the edge server
   - Backup and sync port: The port used by inSync clients for communicating with the edge server. The default port is 6061.
   - Registration Key: The registration key of the edge server.

3. Save your changes.

5.6.6 Configuring the edge server with a storage node

If you install storage nodes that are physically apart from the Master server, you can choose to configure the edge server to authenticate requests to the storage nodes. To allow a storage node to recognize the edge server, you must register the edge server with the storage node. Registration ensures that the storage node recognizes the edge server via which Master server or inSync client requests are routed. Registering the edge server thus ensures that the "handshake" between the storage node and the edge server is complete.

Note: If your Master server and storage node communicate via LAN, configure the same edge server for your Master server and storage node. This is because a single edge server can be configured to a combination of Master server and storage nodes.

Before you begin, make sure that:

- You installed the Master server and the storage node.
- You can log on to the computer on which the edge server in installed.
- Port 6061 on the edge server is free.

Step 1: Generate the edge server key

The registration key is required for establishing a connection between the storage node and edge server. It acts as a unique identifier and ensures that the storage node can identify and interact with the edge server.

Note: Druva recommends that you generate the registration key only once. If you generate the registration key again, the old registration key becomes inactive. Use the key that you generate to register the edge server with existing as well as new storage nodes.

To generate an edge server key

1. Log on to the edge server.
2. Generate the registration key:
   Windows: From the **Start** menu, click **All Programs** > **Druva inSync Edge Server**.

   Linux: Run the following command:
   ```
   sudo insync-edgeserver-config.sh -k
   ```
3. Copy the registration key.

Step 2: Register the edge server with the storage node

Follow the steps in this section to register the edge server with the storage node.

To register the edge server

1. Log on to the storage node.
2. Double-click the Druva inSync Storage Node icon.
3. On the Druva inSync Storage Node window, click **Advanced Options**.
4. Enter the following details:
   - Enable Edge Server: Indicates that edge server association is enabled
   - IP Address/FQDN: The IP address or the fully qualified domain name (FQDN) of the edge server
   - Backup and sync port: The port used to connect with the edge server. The default port is 6061.
   - Registration Key: The registration key of the edge server.
4. Save your changes.

Step 3: (For new storage nodes) Modify the storage node settings

Follow the steps in this section to change the storage node settings from the inSync Master Management Console.

To change to the storage node settings

1. Register the storage node on the Master server.
2. On the second step of the wizard, enter the following details:
   - Enable Edge Server: Indicates that edge server association is enabled
   - IP Address (or) FQDN: The IP address or the fully qualified domain name (FQDN) of the edge server
   - Edge Server Port: The port used to connect with the edge server. The default port is 6061.
3. Click **Create Storage Node**.

(For existing storage nodes) Modify the storage node settings

Follow the steps in this section to change the storage node settings for existing storage nodes.

To change the storage node settings

1. Log on to the inSync Master Management Console.
2. From the menu bar, select **Manage** > **Storage Nodes**> **<storage node name>**.
3. Under Advanced Options, click **Edit**.
4. Enter the following details:

- **Enable Edge Server**: Indicates that edge server association is enabled
- **IP Address (or) FQDN**: The IP address or the fully qualified domain name (FQDN) of the edge server
- **Edge Server Port**: The port used to connect with the edge server. The default port is 6061.

5. Save your changes.

5.7 Performing additional tasks

This section contains a list of miscellaneous tasks that you might want to perform in order to ensure that your customers can start working with their Private Cloud setup. When you created organizations, administrators of the organizations should have received invitation emails. Administrators can log on to the inSync Master Management Console using these credentials.

- To ensure that administrators understand how to use inSync Private Cloud, direct them to the inSync On-premise Help.
- To ensure that end-users of organizations understand how to use inSync client, direct them to the inSync client Help.
- You can choose to download the inSync client installers for your organizations. To do so, click **Download inSync Client** in the right-navigation pane of the Organization Portal.
- You can download the edge server for organizations that choose to set up the edge server on their own. To download the edge server, click **Download Edge Server in the right navigation pane.**