



Cisco Spark Hybrid Directory Service

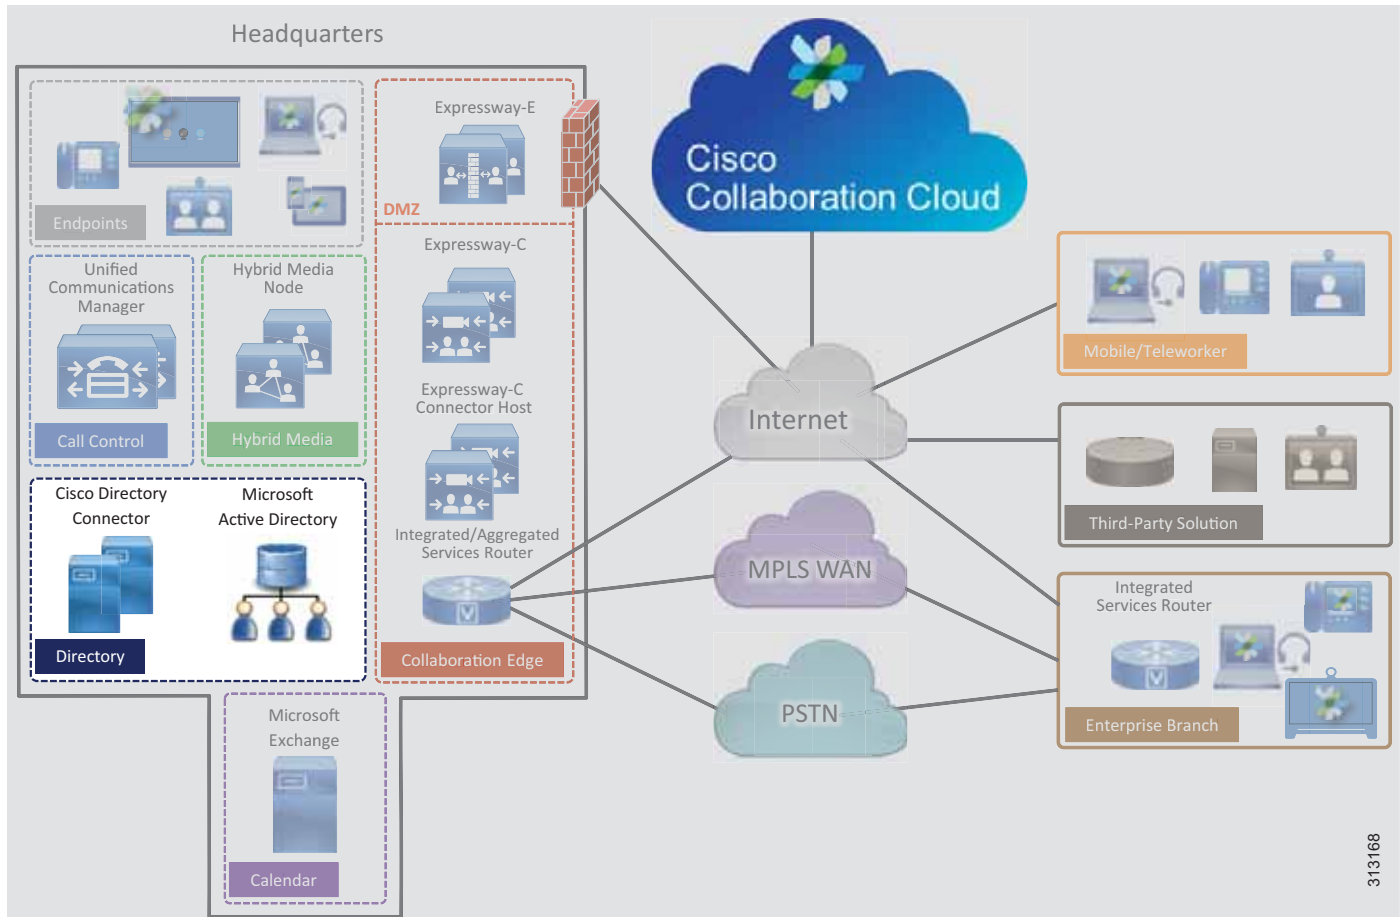
Revised: December 21, 2017

Cisco Spark Hybrid Services enable Cisco Spark customers to connect on-premises collaboration services to the Cisco Collaboration Cloud. Integrating directory services between the on-premises LDAP directory and the common identity service within the customer's Cisco Spark organization adds value by simplifying user on-boarding.

Overview

The Cisco Spark Hybrid Directory Service high-level architecture depicted in [Figure 2-1](#) allows the Cisco Spark customer to synchronize their corporate Microsoft Active Directory with the identity store of their organization in the Cisco Collaboration Cloud. This makes Cisco Collaboration Cloud user on-boarding and service provisioning simple and consistent.

Figure 2-1 Cisco Spark Hybrid Directory Service High-Level Architecture



Prerequisites

Prior to implementing and deploying Cisco Spark Hybrid Directory Service, perform the following requirements:

- Deploy Microsoft Active Directory within the organization and populate it with user information.
- Make sure Cisco Unified Communications Manager (Unified CM) is fully integrated with Microsoft Active Directory (directory synchronization and authentication).
- If the on-premises network is behind a firewall, ensure that outbound access to the Internet through HTTPS on port 443 is available either directly or by way of an HTTP proxy.

Core Components

The core components for Cisco Spark Hybrid Directory Service include:

- Cisco Directory Connector
- Microsoft Active Directory

Recommended Deployment

To deploy Cisco Spark Hybrid Directory Service in the PA for Cisco Spark Hybrid Services, we recommend the following:

- Ensure that the end-user account mail ID field in the Unified CM End User database contains the user's email address. Cisco Spark users correlate to Cisco Unified CM end users by means of email addresses. With LDAP directory integration, the mail ID field for Unified CM end users is typically mapped from the mail field of the LDAP directory during synchronization.
- Install Directory Connector on a separate Windows server from the Active Directory Domain Service or Active Directory Lightweight Directory Services.
- Run a first synchronization after the Directory Connector installation finishes. Then configure full synchronization and incremental synchronization schedules to keep the Directory Connector (and in turn the Cisco Collaboration Cloud) updated when user information changes (user update, deletion, or addition) within Microsoft Active Directory.

Key Benefits

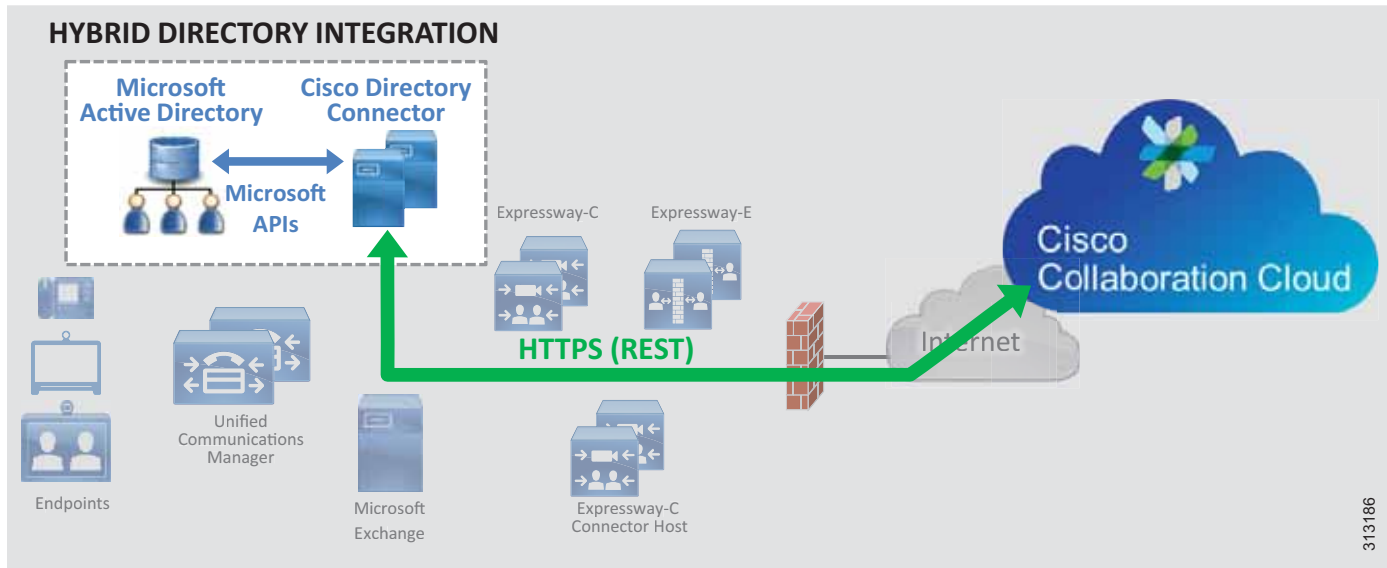
Cisco Spark Hybrid Directory Service provides the following benefits:

- Synchronization of identities, users, and groups, from the corporate Microsoft Active Directory to the cloud, and the creation of Cisco Spark user accounts from this corporate directory source.
- HTTPS outbound connection from the enterprise to the Cisco Collaboration Cloud on standard port 443, which is typically allowed by organizations and thus should not require additional configuration to open ports on the firewall. The organization's existing HTTP proxy may also be leveraged as required.
- Automatic, scheduled synchronization of users from the enterprise Active Directory to the Cisco Collaboration Cloud through the Cisco Directory Connector.
- Incremental synchronization and full synchronization to facilitate management of user identity information.
- Custom attribute mappings between Microsoft Active Directory and Cisco Directory Connector for maximum flexibility.

Architecture

Figure 2-2 shows the Cisco Spark Hybrid Directory Service integration to the enterprise directory. This integration relies on the Cisco Directory Connectors, which are co-located in the central site with the Microsoft Active Directory. Cisco Directory Connector is deployed on two Microsoft Windows Servers for redundancy and high availability.

Figure 2-2 Architecture for Integration of Cisco Spark Hybrid Directory Service with the Enterprise Directory



Cisco Directory Connector relies on Microsoft Active Directory application programming interfaces (APIs) to pull user information from the Microsoft Active Directory. The APIs are based on the Microsoft NET framework. Directory Connector uses HTTPS to push user information to the organization's Cisco Collaboration Cloud common identity store.

Role of Cisco Directory Connector

Cisco Directory Connector plays the role of synchronization agent between the corporate Microsoft Active Directory and the organization's common identity store in the Cisco Collaboration Cloud. Cisco Directory Connector initially populates the Cisco Collaboration Cloud with user information from the Active Directory and maintains this user information with subsequent synchronizations to update the organization's cloud common identity store with the latest moves, adds, changes, and deletions occurring on the enterprise Active Directory.

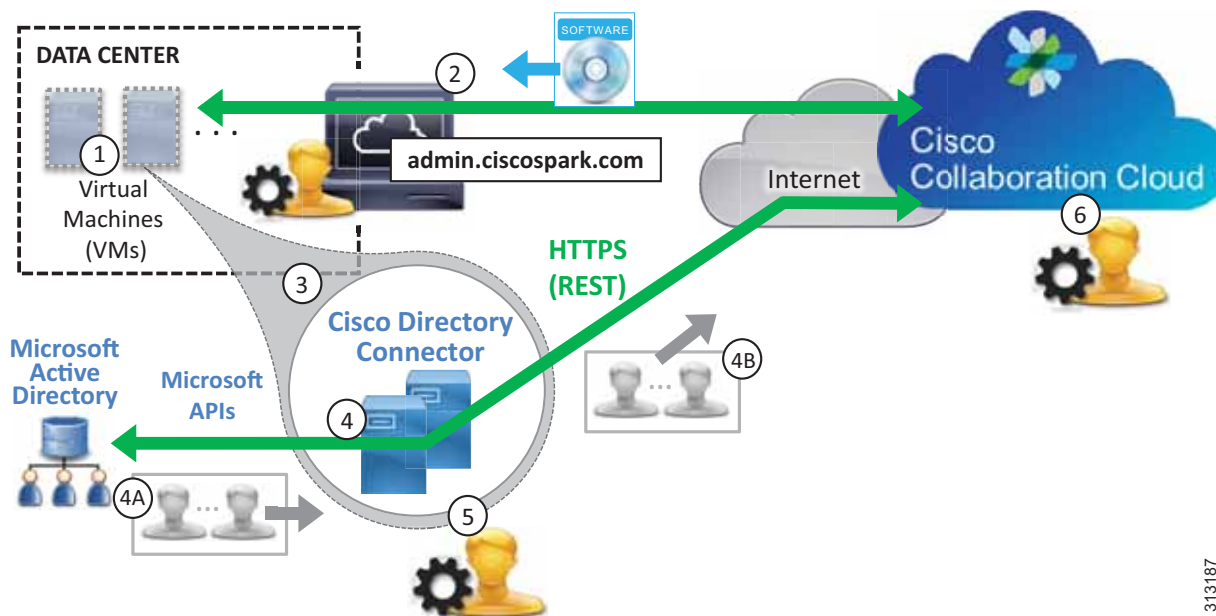
Role of Microsoft Active Directory

Microsoft Active Directory is the enterprise user repository and the single source of validation for user information. The directory administrator maintains the user information contained within the directory with moves, adds, changes, and deletions. Any updates to user information in Active Directory are propagated to the Cisco Directory Connector (and in turn to the Cisco Collaboration Cloud) during synchronization.

Deployment Overview

Figure 2-3 shows the high-level steps required to deploy Cisco Spark Hybrid Directory Service. Virtual Microsoft Windows Servers are created and deployed in the enterprise data center (step 1). After the Windows servers are deployed, the administrator logs into the Cisco Spark Control Hub at <https://admin.ciscospark.com/> to enable directory synchronization and download the Cisco Directory Connector software installation package (step 2). Next, Directory Connector is installed on the Windows servers (step 3). After Directory Connector is installed, the administrator configures the connector (step 4) and an initial synchronization occurs between Microsoft Active Directory and the Directory Connector (step 4A) and between the Directory Connector and the Cisco Collaboration Cloud (step 4B).

Figure 2-3 Cisco Spark Hybrid Directory Service Deployment Overview

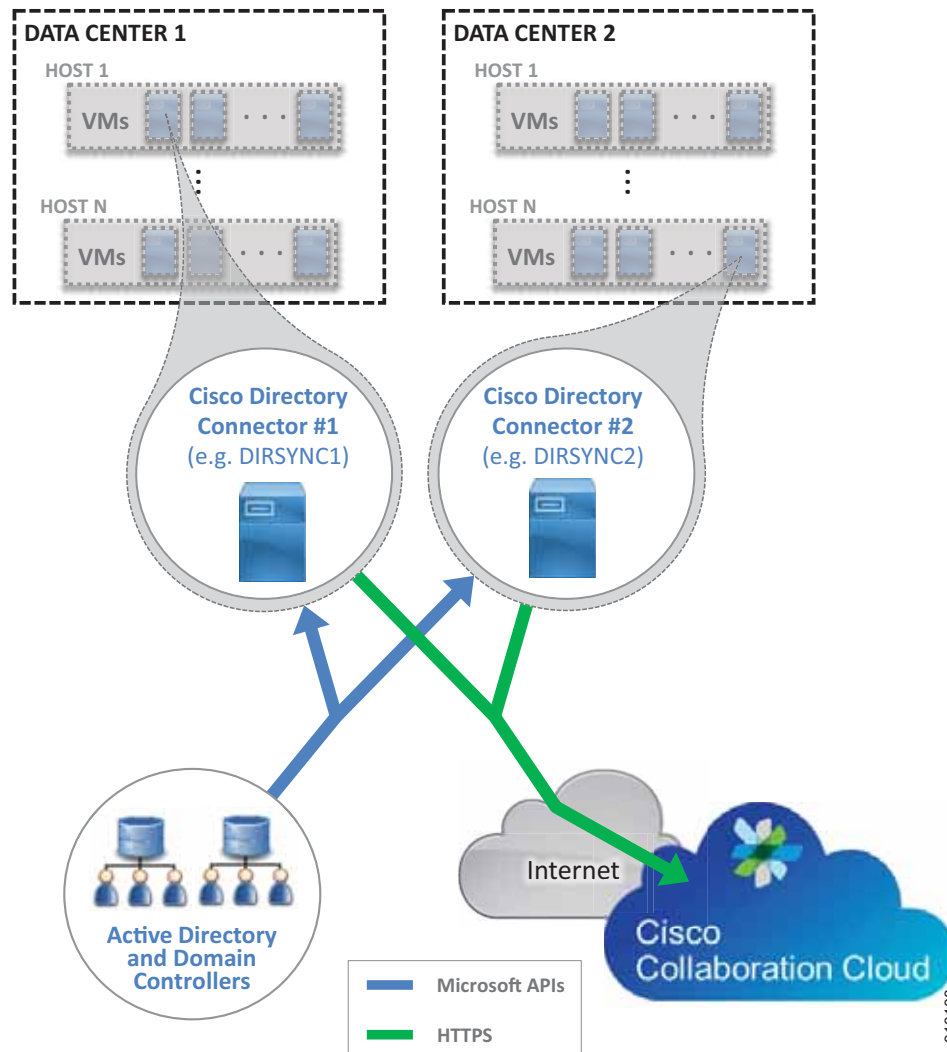


Once the initial synchronization completes, the administrator configures the schedule for periodic incremental and full synchronizations (step 5). After that, the administrator manages users and provisions them for cloud services as appropriate (step 6).

High Availability

As shown in Figure 2-4, two Cisco Directory Connectors are deployed. These Windows Servers virtual machines are deployed on separate hosts in separate buildings or data centers to provide high availability and redundancy. Directory Connectors are deployed as a pair, and both are capable of synchronizing user information between the enterprise directory and the cloud. However, under normal operation, one Directory Connector (primary) handles directory synchronization while the other (backup) maintains connectivity to the Cisco Collaboration Cloud but does not perform any synchronization. In the event that the primary Directory Connector fails, the backup Directory Connector will continue to handle synchronization operations based on the configured failover interval.

Figure 2-4 Cisco Spark Hybrid Directory Service High Availability

**Note**

In cases where only a single Cisco Directory Connector is deployed (non-redundant deployments), if the Directory Connector fails, user information is no longer synchronized between Active Directory and the Cisco Collaboration Cloud common identity store. The administrator is able to manage existing users and to provision them for services while the Directory Connector is down, but no users can be added or removed from the cloud identity store until the Directory Connector is returned to service.

In addition to Cisco Directory Connector high availability considerations, also consider providing redundancy for other aspects of the integration such as the Active Directory services, connectivity to the Cisco Collaboration Cloud (HTTPS), and availability of cloud services.

Microsoft components (Active Directory, Domain Controllers, and other Microsoft enterprise network services) should be deployed in a redundant fashion. Consult Microsoft product documentation for information on high availability.

Highly available network connectivity to the Internet is also required to ensure Cisco Spark and other Cisco Collaboration Cloud services are reachable from the enterprise. Redundant physical Internet connections, preferably from different providers, are recommended.

Cisco Collaboration Cloud services are highly available because those services and components are deployed across multiple physical data centers on elastic compute platforms.

Scalability

The primary sizing and scalability considerations for Cisco Spark Hybrid Directory Service is the size of the synchronization. The larger the enterprise directory and the search base in terms of number of users, the longer a synchronization will take to complete. For this reason it is important to monitor synchronization operations initially to ensure that both incremental and full synchronizations are completing prior to the beginning of the next synchronization period. We recommend running the Directory Connector on a dedicated Windows server host. Additional load on the Windows server can reduce performance and increase overall system response and synchronization times.

For more information on Cisco Spark Hybrid Directory Service scaling, see the chapter on [Sizing Cisco Spark Hybrid Services](#).

Cisco Spark Hybrid Directory Service Deployment Process

Cisco Spark Hybrid Directory Service requires the deployment of the Cisco Directory Connector and synchronization between the on-premises directory and the organizations Cisco Collaboration Cloud common identity store.

Directory synchronization allows corporate users to be imported into Cisco Collaboration Cloud. Directory synchronization is facilitated using Cisco Spark Control Hub and the Cisco Directory Connector. The Directory Connector allows for automatic synchronization of corporate users with the Cisco Collaboration Cloud. Without Directory Connector, users must be imported manually to the Cisco Collaboration Cloud using a **.csv** file.



Note

This section presents high-level guidance for deploying Cisco Spark Hybrid Directory Service. This guidance should be used in conjunction with the detailed instructions provided in the latest version of the *Deployment Guide for Cisco Directory Connector*, available at <https://www.cisco.com/c/en/us/support/unified-communications/spark/products-installation-guides-list.html>.

The deployment of Cisco Spark Hybrid Directory Service starts with the Windows Server installation followed by the download, installation, and initial configuration of Cisco Directory Connector. To deploy Cisco Spark Hybrid Directory Service, perform the following tasks in the order listed here:

1. [Deploy Microsoft Windows Server hosts for Cisco Directory Connector.](#)
2. [Enable directory synchronization and download Cisco Directory Connector software from the Cisco Spark Control Hub.](#)
3. [Install Cisco Directory Connector on the Windows Server host.](#)
4. [Configure Directory Connector and complete the initial synchronization.](#)
5. [Schedule periodic incremental and full synchronizations.](#)
6. [Manage imported users and provision them for Cisco Collaboration Cloud services.](#)

1. Deploy Microsoft Windows Server hosts for Cisco Directory Connector.

The Cisco Directory Connector runs on a trusted Microsoft Windows domain server deployed in the corporate network. The server joins the Active Directory domain and needs an administrator read-only account to authenticate the Cisco Directory Connector server to the on-premises domain.

Deploy a new Microsoft Windows Server and join the corporate Microsoft Active Directory domain. To ensure a highly available deployment of Cisco Spark Hybrid Directory service, install a second domain Microsoft Windows Server on a separate host.

For information about the specific Microsoft Windows Server and Microsoft Active Directory versions supported for Cisco Spark Hybrid Directory Service, refer to the latest version of the *Deployment Guide for Cisco Directory Connector*, available at

<https://www.cisco.com/c/en/us/support/unified-communications/spark/products-installation-guides-list.html>

**Note**

Microsoft Windows Servers should be deployed and configured according to corporate standards and policies and should adhere to any requirements around virus and malware protection, device management, and security.

2. Enable directory synchronization and download Cisco Directory Connector software from the Cisco Spark Control Hub.

Log into the Cisco Spark Control Hub at <https://admin.ciscospark.com/> from the web browser on the Windows Server host you deployed in Step 1. Use your Cisco Spark organization administrator credentials.

On the Cisco Spark Control Hub, enable directory synchronization by navigating to **Users** and clicking **Manage Users**. Next, click **Enable Directory Synchronization** and choose **Next** to continue. Then click the **Download and Install** link to save the Cisco Directory Connector installation **.zip** file (for example, DirectoryConnector.zip) to the local server.

3. Install Cisco Directory Connector on the Windows Server host.

Locate the **.zip** file saved to the host server in Step 2. Unzip the file, navigate to the setup folder, and run the **.msi** file (for example, CiscoDirectoryConnector.msi) in the setup folder to launch the Cisco Directory Connector Setup wizard.

Select **I accept the terms in the License Agreement** and click **Next** to accept the license agreement. Click **Next** to accept the default installation location.

Select the **Domain Account** option for the service account and enter the username and password for the domain account. In the **Username** field include the Active Directory domain and username with the format `<domain>\<user_name>` (for example, ENT-PA\administrator). Click **Next** to save the domain account information.

Then click **Install** to start the installation of Cisco Directory Connector.

When the installation completes, repeat this step on the second Windows Server host to install a redundant Directory Connector.

4. Configure Directory Connector and complete the initial synchronization.

Launch Cisco Directory Connector and sign into the Cisco Spark organization by entering the email address and password of the administrator account for the organization. Note that this is the same email address and password used to log into the Cisco Spark Control Hub management portal. Click to confirm the Cisco Spark organization and domain.

Next, perform initial configuration of Directory Connector. From the Directory Connector dashboard click the **Configuration** tab.



Note

If a configuration tab or field value is not mentioned here, then the default setting and value should be assumed.

Navigate the tabs on the Configuration screen and configure the settings as shown in [Table 2-1](#).

Table 2-1 Cisco Directory Connector Configuration Settings

Configuration Tab	Setting	Description and Values
General	Connector Name	Enter a name for the Directory Connector (for example, DIRSYNC1). This is the name that will be displayed on the dashboard and on the Cisco Spark Control Hub web portal.
	Preferred Domain Controllers	Add one or more Domain Controllers by using the drop-down menu. Select a Domain Controller on the network and click the Add button. Add at least two Domain Controllers to ensure directory services are highly available on the network.
Object Selection During synchronization, user information (based on the selected containers and configured LDAP filters) is pushed from the Directory Connector to the Cisco Collaboration Cloud through an HTTPS connection. Because this is an outbound connection from the enterprise's perspective, it does not require any inbound ports to be opened on the internal or external firewall.	Object Type	Users box is selected.
	LDAP Filters	Enter any required LDAP filters in standard LDAP format to limit the number of searchable containers. Directory Connector pulls user information from the corporate Microsoft Active Directory. User information can be pulled from the entire domain or from specific containers and organizational units. Create multiple LDAP filters if more granularity is needed.
	On Premises Base DN's to Synchronize	Select one or more DN's from the window (for example, CN=Users,DC=ent-pa,DC=com) and click Select to choose appropriate synchronization containers and objects (for example, Users) to include in the directory synchronization agreement. The Cisco Spark Hybrid Directory service integration with Microsoft Active Directory supports deployments with both single and multiple forests and either single or multiple domains. Multiple domains are supported when Active Directory Lightweight Directory Services (AD LDS) is deployed.
User Attribute Mapping Cisco Directory Connector synchronizes a number of Microsoft Active Directory attributes to the Cisco Collaboration Cloud (identity store of the customer organization).	Active Directory Attribute Name	Configure any required Active Directory to Cisco Collaboration Cloud attribute name mappings by selecting options from the Active Directory attribute drop-down lists. At a minimum, ensure that the Active Directory attribute name mail is mapped to the required Cisco Collaboration Cloud attribute name uid . The mail attribute plays a key role in the Cisco Collaboration Cloud because it uniquely identifies the user. Other commonly synchronized Microsoft Active Directory attribute names include displayName, givenName, and telephoneNumber.

Click **Apply** to save and apply the configuration settings.

Once Directory Connector is installed and configured as above, perform an initial full synchronization to pull user information from the corporate Microsoft Active Directory and push it to the organization's Cisco Collaboration Cloud common identity store.

On the redundant Cisco Directory Connector, configure the same settings shown in [Table 2-1](#), but for the Connector Name setting use a unique name (for example, DIRSYNC2).

5. Schedule periodic incremental and full synchronizations.

After the initial synchronization, it is important to keep the organization's Cisco Collaboration Cloud common identity store updated with moves, adds, and changes that occur in the corporate Active Directory.

To keep the Cisco Collaboration Cloud up to date with corporate directory changes, configure periodic incremental and full synchronizations on one of the Directory Connectors. Return to the Directory Connector Configuration tab and select **Schedule**. Then configure synchronization settings as shown in Table 2-2.

Table 2-2 Cisco Directory Connector Schedule Configuration Settings

Configuration Tab	Setting	Description and Values
Schedule	Incremental Sync Interval	Set the interval in minutes between incremental synchronizations (for example, 10 minutes is the default).
	Enable Full Sync Schedule	Select this option.
	Schedule	Select the time and day(s) of week to perform a periodic full synchronization (for example, 7:30 AM on S(unday)).
	Failover Interval	Set the amount of time in minutes before the secondary Directory Connector becomes primary and takes over incremental and full synchronization (for example, 60 minutes is the default). This setting applies for high availability deployments with more than one Directory Connector.

The settings in Table 2-2 are shared and apply to both Directory Connectors in the deployment.

6. Manage imported users and provision them for Cisco Collaboration Cloud services.

After the enterprise directory user information has been propagated to the Cisco Collaboration Cloud, the administrator is able to provision users for cloud services and manage those service features and settings by using the Cisco Spark Control Hub.

Use your Cisco Spark organization administrator credentials to log into the Cisco Spark Control Hub at <https://admin.ciscospark.com/> from a web browser.

On the Cisco Spark Control Hub, begin managing and provisioning user services by navigating to **Users** and then clicking **Manage Users**. Once directory synchronization is enabled, there are multiple ways to modify users and the services they use. Users can be modified manually or in bulk.

To modify large numbers of users in bulk, choose either **Export and modify users with a CSV file** or **Modify all synchronized users**. The CSV file method is good for modifying groups of users in bulk (up to 1,100 users at a time); however, preparing the CSV file for bulk modification is a manual process.

To enable a feature or service for all users, click **Modify all synchronized users** and click **Next**. On the next screen, wait for the system to synchronize the list of users from the latest synchronization agreement and then click **Next**.

On the subsequent screen, provision users for Message, Meeting, and other services including Hybrid Services. Once you have selected the services, click **Next** to start the update of user accounts. When the update is complete, users can begin to use the added services and features.

**Note**

Valid licenses are required to add and enable licensed services and features.
