IBM Cloud

Cloud Shell

Product guide

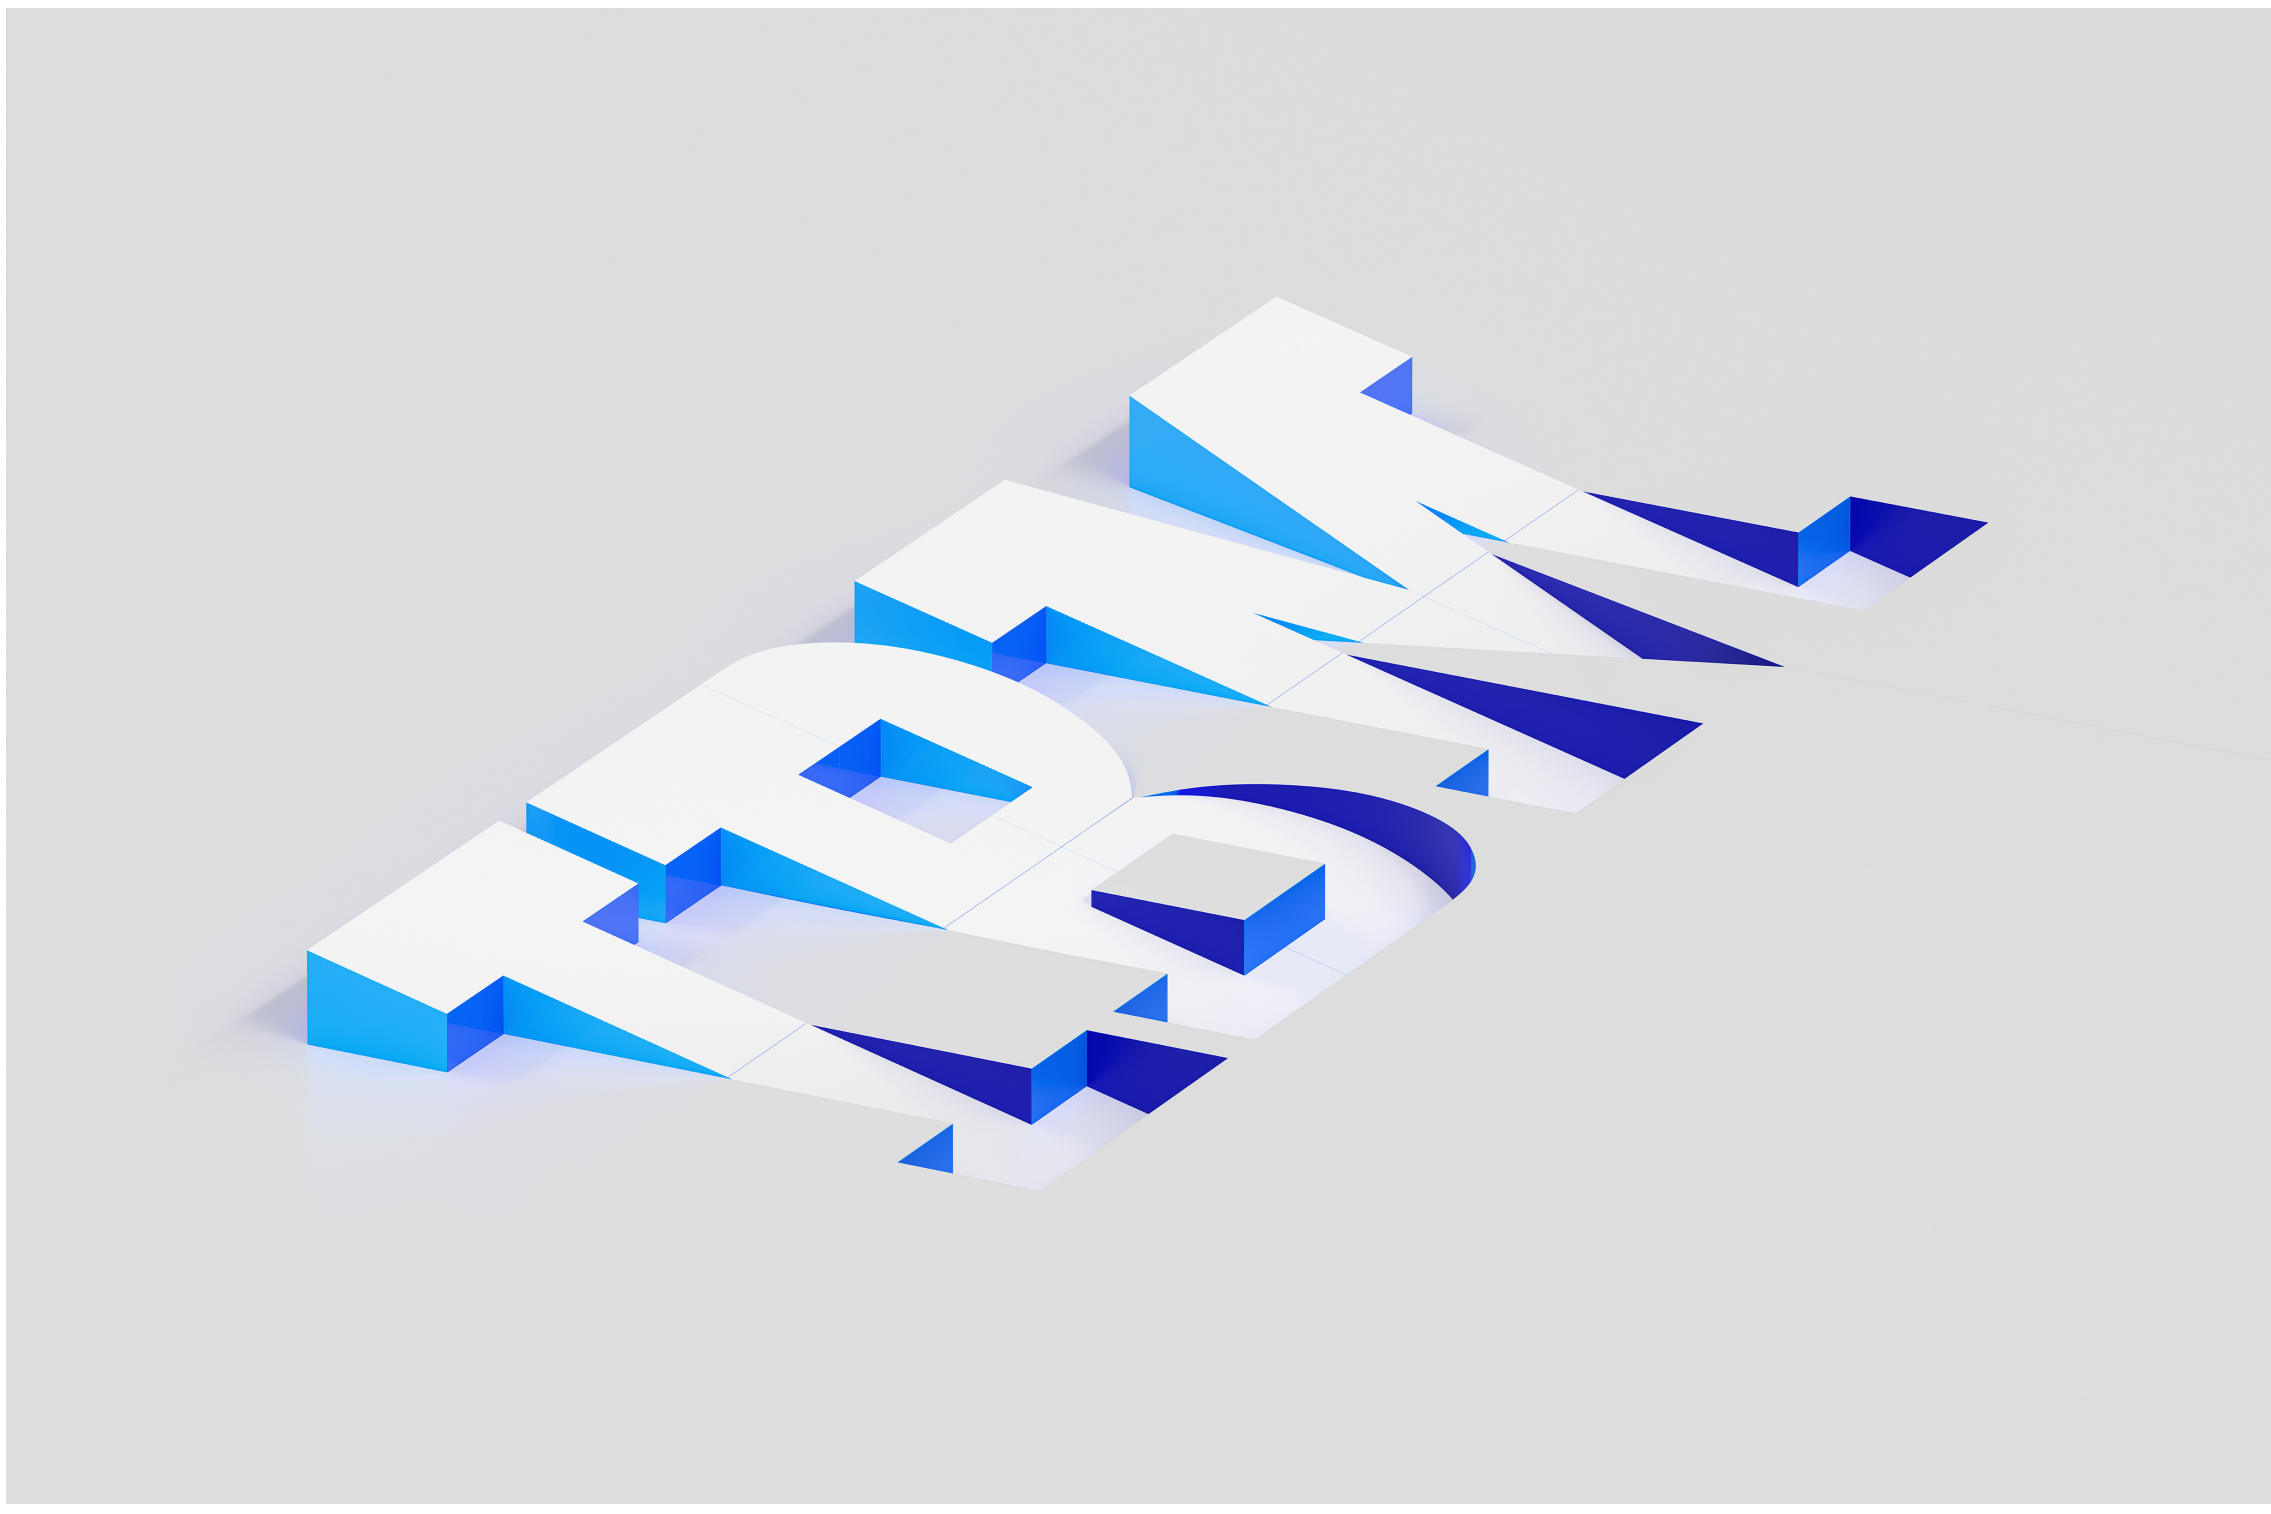

IBM

Edition notices

This PDF was created on 2024-12-13 as a supplement to *Cloud Shell* in the IBM Cloud docs. It might not be a complete set of information or the latest version. For the latest information, see the IBM Cloud documentation at https://cloud.ibm.com/docs/cloud-shell.

Getting started with IBM Cloud Shell

In this getting started tutorial, you use IBM® Cloud Shell to clone a sample Node.js app. IBM Cloud Shell is a cloud-based shell workspace that you can access through your browser. Cloud Shell is preconfigured with the full IBM Cloud CLI and tons of plug-ins and tools that you can use to manage apps, resources, and infrastructure.

Tip: New to Cloud Shell? Watch the demo video to see how you can start to use the service from the IBM Cloud console.

Before you begin

- If you don't already have an IBM Cloud account, sign up for one and log in.
- Note the following account access information that might affect your access to Cloud Shell and its features:
  - An account administrator can grant specific users access to Cloud Shell and its features, even if Cloud Shell settings are disabled at the account level. For more information, see Assigning access to Cloud Shell and its features at a user level.
  - Account owners or users with Cloud Shell administrator access can enable or disable Cloud Shell features for an account. The available features in this release are **File upload and download** and **Web preview**. The feature settings apply only to the enabled Cloud Shell locations. For more information, see Enabling or disabling Cloud Shell features for an account.
  - For more information, see IAM roles and actions.

Step 1: Start a session in IBM Cloud Shell

In the IBM Cloud console, click the IBM Cloud Shell icon . A session starts and automatically logs you in through the IBM Cloud CLI.

Step 2: Clone the Node.js sample app

Your Cloud Shell session starts in the temporary home directory, `/home/<user-name>`. Clone the Node.js Express sample app to this directory and then change to the sample app directory by running the following commands.

```
$ git clone https://github.com/IBM/nodejs-express-app.git
```

```
$ cd nodejs-express-app
```

From your app directory, run the following command to view the new files.

```
$ ls
```

You can see that your `nodejs-express-app` directory contains all of the sample app files and folders.

```
Dockerfile  Dockerfile-tools  LICENSE  README.md  package-lock.json
package.json  public  scripts  server  test
```

Step 3: Preview the app locally

You can start it in Cloud Shell and preview it locally.

1. First, download project dependencies for the Node.js sample app.

   ```
   $ npm install --only=prod
   ```

2. Start the app.

   ```
   $ npm run start
   ```

   When the app is running, it outputs the location where the app UI is available. For Cloud Shell, you need to note only the port, which in this case is `3000` .

   ```
   > nodejsexpressapp@1.0.0 start /home/my-user-name/nodejs-express-app
   > node server/server.js

   App UI available http://localhost:3000
   ```

```
Swagger UI available http://localhost:3000/swagger/api-docs
```

3. In the Cloud Shell menu bar, click the **Web preview** icon  , and select port 3000.

   Your app preview opens in a new window. If you see the congratulations message, your app is running locally in Cloud Shell!

4. Press `Control+C` to stop the app.

Next steps

From Cloud Shell, you can quickly run commands on your IBM Cloud account. For example:

- Managing IAM access, API keys, service IDs, and access groups
- Viewing billing and usage information

In addition to the IBM Cloud CLI, Cloud Shell is preconfigured with all IBM Cloud CLI plug-ins and many tools and runtimes to help you work in the cloud. For the full list, see Installed plug-ins and tools.

Working in Cloud Shell

IBM® Cloud Shell contains a personal workspace and sessions where you can run commands. You can open up to five concurrent sessions, which operate independently so you can work with different resources, regions, and accounts at once.

To open Cloud Shell, click the IBM Cloud Shell icon in the IBM Cloud console. A session starts and automatically logs you in to the IBM Cloud CLI with your current account.

Your Cloud Shell workspace

Each time that you open Cloud Shell, you connect to your Cloud Shell workspace. A workspace is a dedicated environment that holds your user and session data and provides temporary storage for you to work with files. Workspaces are separate for each user, account, and region, which each run in their own virtual machine. Only you can access your workspace, and if you're a user in multiple accounts, you have separate workspaces for each account, which keeps your data that is isolated and secure. For more information about storing data in Cloud Shell, see Temporary workspace storage.

Deleting your Cloud Shell workspace

Deleting your Cloud Shell workspace ends and deletes all sessions in the workspace. All files and data within the workspace are erased. Before you delete the workspace, download any files that you want to keep.

To delete your Cloud Shell workspace, complete the following steps:

1. Download any files that you want to keep.
2. From your Cloud Shell workspace, click the **Actions** icon , and then select **Delete**.
3. In the confirmation window, click **Delete**. A blank page is displayed.
4. Click the **Add** icon + to open a new Cloud Shell session. This action also creates a new Cloud Shell workspace.

Working with sessions

A Cloud Shell session is the terminal interface where you can run commands, scripts, and other tools. It's similar to the command prompt on your local machine, but it's preconfigured with the full IBM Cloud CLI and numerous plug-ins and tools so that you can work in IBM Cloud without having to install anything. For more information about the CLI, see Getting started with the IBM Cloud CLI.

When you open a session, you start in your Cloud Shell workspace's home directory, `/home/<user-name>`. Your home directory provides a temporary place for you to work with files in Cloud Shell. It is shared between all of your sessions. Because data in your home directory isn't persisted, any files that you add or change are removed after two hours of inactivity. For more information about your Cloud Shell workspace and home directory, see Working with files.

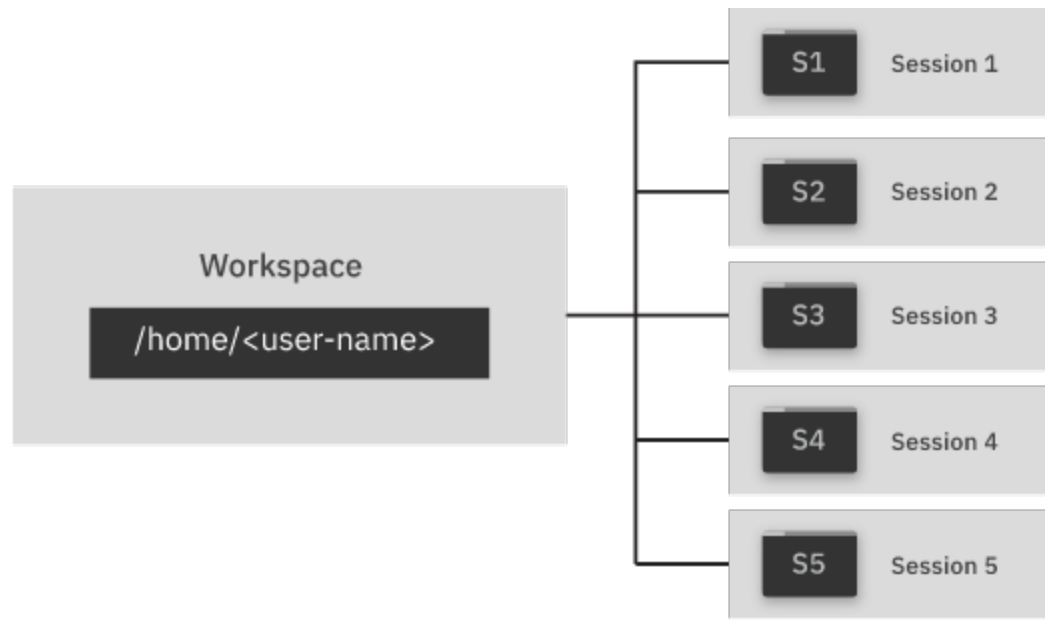

Your personal workspace and sessions in Cloud Shell

Opening sessions

To open a new session, click the Open icon +. You can have up to five sessions open at one time. Each session maintains separate command histories, which are removed when you close the session. Because the workspace storage is shared between sessions, you can save a file to your home directory in one session, and then open the file in another session.

Tip: Each session automatically closes if it's not used for an hour. If you don't use Cloud Shell for two hours - that is, another hour after the last session is closed - all files and data in your Cloud Shell workspace are erased. If you need to step away from Cloud Shell, be sure to download any files that you want to keep. For more information, see Working with files.

Switching accounts

When you open a session, you're logged in with the IBM Cloud CLI targeting the account and region that you were in when you opened Cloud Shell. This account and region are the default for every session.

To switch the default account for your sessions, close Cloud Shell, switch to a different account in the IBM Cloud console menu bar, and then reopen Cloud Shell.

Because you have separate workspaces for each account, your session history and temporary file storage are also separate. This means that when you switch accounts, you see different session tabs and don't have access to any files that you were working with in the previous account.

Changing locations

Cloud Shell runs in the Dallas (`us-south`) and Frankfurt (`eu-de`) regions. Each Cloud Shell location operates independently, which means that your workspace storage, sessions, and usage are separate. When you work in Cloud Shell, the location that you're in determines the following aspects:

- Where your workspace data is stored, including any files that you upload and your user data, such as your session history and usage quota
- The default region that's targeted by the IBM Cloud CLI: `us-south` or `eu-de`

By default, the location that's closest to you geographically is selected when you open Cloud Shell from the console. For example, if you're in California then Cloud Shell opens to the Dallas region, but if you're in Paris it opens to the Frankfurt region.

To change to Cloud Shell in a different location, find the current location in the Cloud Shell menu bar, and click **Change**. Select the location that you want to switch to, and click **OK**.

If you change from the default location, Cloud Shell remembers your location preferences by saving them as cookies in your browser. If you switch to a different browser, clear your browser cookies, or your cookie preferences don't allow personalization cookies, Cloud Shell opens in the default location.

Restarting Cloud Shell

When you restart Cloud Shell, any open sessions are closed and all of your files and data are removed. Before you restart, be sure that you download any files that you want to keep to your local system.

To restart, click the Menu icon , and select **Restart**.

Tracking your usage

You can use Cloud Shell for up to 50 hours within a week per region. Anytime that you have at least one open session counts toward this usage quota, even if you aren't actively running commands. For example, if you're working in Cloud Shell for 15 minutes and then leave your session open for a 30-minute lunch break, those 45 minutes are counted as usage. If you have three sessions open at the same time over an hour period, your usage is one hour because concurrent sessions don't count as additional usage. To minimize your usage, be sure to close sessions after you're done using them.

If you use all of your quota, your Cloud Shell sessions are closed. Closing Cloud Shell removes any data in your workspace, including your files and command history.

Tip: When you reach your usage quota, you get a notification that Cloud Shell will close in 5 minutes. Use this time to complete any urgent tasks.

At any point, you can track your usage and check when your weekly quota resets. In the Cloud Shell menu bar, click the Menu icon , and select **Usage quota**. You can view how many hours you used Cloud Shell during the current weekly period and the day and time when it resets.

The time that you used of the 50-hour usage quota resets back to 0 at the same time each week. Even after you used all of your quota, you can still use IBM Cloud Shell for up to 5 minutes at a time to complete quick tasks, such as to run a few simple commands.

Note: IBM Cloud Shell is intended for IBM Cloud administration and development purposes. Any usage outside of these guidelines might result in your access to IBM Cloud Shell or IBM Cloud being suspended. Because Cloud Shell is intended for interactive purposes only, long-running processes are stopped. For more information, see the IBM Cloud Terms and Notices.

Changing your cookie preferences

In Cloud Shell, browser cookies are used to provide a secure login and remember some preferences, such as if you change the default location. Additional cookies are used to build a more personalized experience for you and provide better support.

Cloud Shell uses the cookie preferences from the IBM Cloud console. If you change your cookie settings from the IBM standard default, your Cloud Shell

location preferences aren't saved, and you might not be able to use some IBM Cloud console features.

To change your IBM Cloud cookie preferences from within Cloud Shell, click the Menu icon , and select **Privacy**. For more information, see Changing your cookie settings.

Working with files

Your IBM® Cloud Shell workspace includes a home directory where you can temporarily work with files in your Cloud Shell sessions. You can upload or download files one at a time to your workspace through the UI, or use command-line tools to work with many files at once.

Before you begin

Enable **File upload and download** in the IBM Cloud console. For more information about how to enable file upload and download, see Enabling or disabling Cloud Shell features for an account.

Temporary workspace storage

Your Cloud Shell workspace includes 500 MB of temporary secure storage, which you can access through your personal home directory, `/home/<user-name>`. Only you can access your workspace storage, and your storage is separate for each account that you access or region where you work in IBM Cloud Shell.

Your workspace storage is shared across all of your sessions, so you can work with the same files in different session tabs. The storage persists only while your workspace is active. If you're idle in Cloud Shell for over an hour, your files and data are removed. Similarly, if you reach the usage quota or you restart Cloud Shell, Cloud Shell closes and removes your data. If you change files that you want to keep, be sure to download the files at the end of your session.

⚠ **Important:** A known issue exists where your connection to Cloud Shell is lost if your reach the temporary storage limit. If this issue happens, the only way to fix the connection is to restart Cloud Shell, which removes all of your files. While you work in IBM Cloud Shell, avoid uploading large files and remove any unused files by using standard Linux™ commands, such as `rm`.

Backing up your data

Your workspace storage is temporary, and it is not intended to be used as the primary storage location for your files. Don't store business-critical or sensitive data in your workspace, and create a backup of your files outside of IBM Cloud Shell. You're responsible for your data, and your backup can help you recover in case an outage or data loss occurs for any reason. For more information about how to transfer files to create a backup, see Transferring multiple files.

Uploading files

You can upload a file to the root level of your home directory. Only a single file can be uploaded at a time.

1. In the IBM Cloud Shell menu bar, click the Upload icon .
2. Select the file that you want to upload, and click **Open**.

Your file is uploaded in your home directory, for example `/home/<user-name>/myFile.txt`. You can move or otherwise work with your files by running standard Linux™ commands. For example, you can move a `myFile.txt` file that you uploaded to a `myFolder` subdirectory by running the following command.

```
$ mv myFile.txt ./myFolder/
```

Tip: Although you can move files, be sure to keep all files in your workspace under your home directory, `/home/<user-name>`. If you move files outside of this directory, it can cause Cloud Shell to close, which removes your data.

Downloading files

You can download a file from your workspace to your local system. Only a single file can be downloaded at a time.

1. Find the path to the file from the command line by using standard Linux commands.

   For example, you can list all files and subdirectories within your current directory.

   ```
   $ ls -R
   ```

   Or, you can search for a file name. The following command searches for files with `myFile` in the name.

   ```
   $ find -iname *myFile*
   ```

2. In the IBM Cloud Shell menu bar, click the Download icon .
3. Enter the path to the file in your home directory, such as `/myFolder/myFile.txt`. Click **Continue**.

> **Tip:** Don't include the home directory root `/home/<user-name>` in the file path. File paths are case-sensitive.

4. Follow your browser prompt to open or save the file to your computer.

Transferring multiple files

Using the Cloud Shell UI, you can upload or download only a single file at a time. If you need to transfer many files, it might take a long time to move all of them individually. Instead, use these strategies to efficiently move files between your workspace and another file system.

Create archives to move files

Before you move files, combine the files into an archive file such as a `.tar`, `.tar.gz`, or `.zip` so that you can move them all at once.

For example, to upload a folder of IBM Cloud administration scripts, you might compress them into a `myScripts.zip` file and upload them to Cloud Shell. In your Cloud Shell session, you can then run `unzip myScripts.zip` to extract the files.

You can do the same thing in reverse when you want to download files. For example, say you want to back up your entire Cloud Shell workspace. From your home directory, run `tar -cvf myTar.tar *` to create a `.tar` file (Mac or Linux) or run `zip -r myZip.zip *` to create a `.zip` file (Mac or Windows). Then, download the archive file from Cloud Shell and extract it on your local system.

Work from a Git repo

For projects in Git repositories, use the Cloud Shell UI to upload an SSH key so that you can connect to your Git repo. Keep a local copy of your SSH key as a backup. Then, run `git clone` to clone all of the repo's files to your Cloud Shell workspace. As a bonus, when you commit and push your changes to your Git repo, your changes are automatically backed up to a file system that's outside of Cloud Shell.

Previewing apps

You can preview apps that are running in IBM® Cloud Shell just like you do if you are running them locally. The Cloud Shell web preview makes your app UI or endpoint available at a URL that only you can access. With this feature, you have a chance to view changes to your app before you deploy it to a publicly accessible location.

Before you begin

Enable **Web preview** in the IBM Cloud console. For more information about how to enable web preview, see Enabling or disabling Cloud Shell features for an account.

Web preview ports

You can preview an app by clicking the **Web preview** icon in the Cloud Shell menu bar. Web preview is available for any app that listens to HTTP requests on the following ports:

- 3000
- 8080
- 8081
- 8082
- 8083
- 8084
- 9080

As with your local system, only one app can run on each port at any time. If you have multiple apps that are listening on the same port, it results in port conflicts.

Previewing apps running directly in Cloud Shell

Cloud Shell includes lots of plug-ins, tools, and runtimes that you can use to work with apps that you're developing. Because Cloud Shell storage is temporary, the easiest way to work with an app is to first copy or clone it to your Cloud Shell workspace, make any changes, and then preview it.

The following example workflow shows how you can preview a Node.js Express sample app that's hosted in a GitHub repository.

1. In a Cloud Shell session, clone your app source to your home directory.

   ```
   $ git clone https://github.com/IBM/nodejs-express-app.git
   ```

   Then, change to the directory where your app was cloned.

   ```
   $ cd nodejs-express-app
   ```

   At this point, you can make any changes to your app code, such as by using the `vim` text editor.

2. Start your app in Cloud Shell.

   To start this Node.js sample app, you first need to download any project dependencies.

   ```
   $ npm install
   ```

   Then, you can start the app so that it's running on a web server in Cloud Shell.

   ```
   $ npm run start
   ```

   After the app starts, note the port where your web server is listening. The port must be one of the available ports to be able to preview it in Cloud Shell. When this sample app starts, it outputs the following information, which shows that it's listening on port `3000`.

   ```
   > nodejsexpressapp@1.0.0 start /home/my-user-name/nodejs-express-app
   > node server/server.js

   App UI available http://localhost:3000
   Swagger UI available http://localhost:3000/swagger/api-docs
   ```

3. In the Cloud Shell menu bar, click the **Web preview** icon  , and select the port where your server is listening. For the Node.js sample app, select 3000.

   Your web preview opens in a new browser window and shows the equivalent of `http://localhost:3000`. You should see the first screen of your app's UI.

   > **Tip:** If you see a message that says there's nothing to preview in Cloud Shell, the server where your app is running isn't listening on the selected port. Verify that your server is listening on the port that you're previewing, and refresh the preview page.

Now that you're previewing your app, you can see any changes that you make by restarting the app if needed and refreshing the page. Be sure to commit any changes back to your GitHub repo or otherwise download your changes because your workspace storage is temporary.

Previewing apps that run in a Kubernetes pod

You can also use Cloud Shell to preview apps that are running in Kubernetes, such as a Prometheus dashboard or your typical web app. Because apps that are running in a pod in a Kubernetes cluster use an internal port, you need to expose the pod's port. You can do that either by creating a service or setting up port forwarding.

If you create a service, you need to expose the service publicly, set up authentication to secure access, and set up an application load balancer (ALB) for the service. For more information, see Connecting applications with services.

If you're doing app development and testing, port forwarding is a simpler way to be able to quickly preview apps in your cluster. It requires much less up-front configuration, but you also must reconfigure it whenever the pod restarts. The following steps walk you through setting up port forwarding and previewing an app that's running in a IBM Cloud Kubernetes Service cluster.

1. In a Cloud Shell session, set the cluster where your app is running as the context. With this setting you can use `kubectl` commands to work with your cluster. For more information, see Installing the CLI.

2. Find the name of the pod where your app is running.

   ```
   $ kubectl get pods
   ```

   The pod name is listed under the `NAME` column, for example `my-pod-123abc`.

   ```
   NAME            READY   STATUS    RESTARTS   AGE
   my-pod-123abc   1/1     Running   0          9d
   ```

3. Find the port in the pod that you want to display by running the `kubectl describe` command on the pod.

   ```
   $ kubectl describe pod my-pod-123abc
   ```

   Under the `Containers` section, the local port is listed as the `Port` value, such as `6000` in the following example.

   ```
   Containers:
   cluster-agent:
   Container ID:   containerd://1234abcd5678efab90cde603e4ba6986fdaf26daae94c4f309
   Image:          us.icr.io/my-registry/my-image:0.5.0-dev.0.1234
   Image ID:       us.icr.io/my-registry/my-image@sha256:abc123def456abc789def012
   Port:           6000/TCP
   Host Port:      0/TCP
   ```

4. Set up port forwarding by running the `kubectl port-forward` command. Be sure to select one of the available ports for web preview in Cloud Shell.

   For example, the following command listens on port `8080` locally and forwards it to port `6000` within the pod.

   ```
   $ kubectl port-forward pod/my-pod-123abc 8080:6000
   ```

   Now that port forwarding is set up, you can preview your app.

5. In the Cloud Shell menu bar, click the **Web preview** icon  , and select the port that's being forwarded to your pod.

   Your web preview opens in a new browser window. You can see the first screen of your app's UI.

   > **Tip:** If you see a message that says there is nothing to preview in Cloud Shell, the server where your app is running isn't listening on the

selected port. Verify that your server is listening on the pod's internal port and that you're forwarding the port that you're previewing to the correct internal port.

Previewing other endpoints on a port

After you open the web preview, you can view any other endpoints that your server is listening for on the same port. The web preview opens at a URL that's the equivalent to `http://localhost:<port>`. To view another endpoint, add the path to the web preview URL.

For example, the Node.js sample app outputs the following URLs:

```
App UI available http://localhost:3000
Swagger UI available http://localhost:3000/swagger/api-docs
```

To view the Swagger UI, add `/swagger/api-docs` to the web preview URL:

`https://p3000-abcd1234-ef56-a1b2-b122-a122560ba24f.dp3.us-south.shell.cloud.ibm.com/swagger/api-docs`

Submitting feedback

Help us improve IBM® Cloud Shell by sharing your experience with our team. Whether you have questions, concerns, or want to give us a rave review, we want to hear from you!

Submit a comment by using the Feedback button

Let our team know your thoughts about Cloud Shell or the documentation by clicking the **Feedback** button on the edge of the page. You can submit feedback about the entire experience or a specific part of the page you're currently on. Your feedback goes directly to our development team. For more information, see Submitting feedback.

Chat with us on Slack

Have a casual question or want to talk about all things IBM Cloud Shell? Come chat with the development team on IBM Cloud Dev Tools Slack. After you request your invitation, sign in and join the `#ask-your-question` channel.

Observability

Auditing events for IBM Cloud Shell

As a security officer, auditor, or manager, you can use the IBM Cloud Activity Tracker service to track how users interact with IBM® Cloud Shell. Cloud Shell automatically generates events that you can analyze in the Activity Tracker service.

Activity Tracker records user-initiated activities that change the state of a service in IBM Cloud. You can use this service to investigate abnormal activity and critical actions and to comply with regulatory audit requirements. In addition, you can be alerted about actions as they happen. The events that are collected comply with the Cloud Auditing Data Federation (CADF) standard. For more information, see IBM Cloud Activity Tracker.

List of events

The following table lists actions in IBM Cloud Shell that generate an event.

Action	Description
`cloudshell.server.create`	An event is generated when a new session is created.
`cloudshell.server.configure`	An event is generated when a session is configured. This event is generated for configuring new sessions and reconfiguring an existing session.
`cloudshell.server.delete`	An event is generated when a session is deleted.
`cloudshell.account-settings.update`	An event is generated when Cloud Shell settings are updated for an account.

Actions that generate events

Viewing events

Events that Cloud Shell generates are automatically forwarded to the Activity Tracker service instance that is available in the same location.

Activity Tracker can have only one instance per location. To view events, you must access the web UI of the Activity Tracker service in the same location where your service instance is available. For more information, see Launching the web UI through the IBM Cloud UI .

Analyzing events

Activity Tracker events contain fields that describe the action that occurred. Values in the `requestData` and `responseData` fields are specific to IBM Cloud Shell, and the other fields are common to all Activity Tracker events. For a more information about common fields, see Event fields.

When a user opens, configures, or closes a Cloud Shell session, the event that is triggered has an `action` field set to `cloudshell.server.create`, `cloudshell.server.configure`, or `cloudshell.server.delete`. The event includes the following fields:

- The `initiator.name` field includes information about the user who interacted with the session.
- The `initiator.id` field shows the IBMid of the user who interacted with the session.
- The `target.id` field includes the Cloud Resource Name (CRN) of the Cloud Shell account and server where the session was modified, in the format `crn:v1:bluemix:public:cloudshell:<REGION>:a/<ACCOUNT_ID>:<CLOUD_SHELL_SERVER_ID>::`

When a session is successfully opened, configured, or closed, the corresponding event that is generated has an `outcome` that is set to `success` with a 200 `reason.reasonCode`. Otherwise, the event has an `outcome` of `failure` with the appropriate HTTP status code in `reason.reasonCode`, and the `responseData` field contains details about the error.

When an account owner or Cloud Shell administrator updates the Cloud Shell settings, the event that is triggered has an `action` field set to `cloudshell.account-settings.update`.

Understanding high availability and disaster recovery for Cloud Shell

As an IBM Cloud® platform service, IBM® Cloud Shell follows the practices that are described in How IBM Cloud ensures high availability and disaster recovery. One exception is that IBM Cloud Shell is regionally available rather than globally available. Because Cloud Shell is region-dependent, automated global failover isn't supported. Always maintain a backup of your data outside of Cloud Shell.

IBM Cloud Shell is available in the Dallas (`us-south`) and Frankfurt (`eu-de`) regions, and each location has three different data centers for redundancy. The data for each location is kept in the data centers near that location. If all the data centers in a location or region fail, IBM Cloud Shell becomes unavailable in that location or region. If this issue happens, you can change to another location or region as described in Changing locations.

For more information, see How IBM Cloud ensures high availability and disaster recovery .

Backing up your data

Your Cloud Shell workspace storage is temporary, and it is not to be used as the primary storage location for your files. Don't store business-critical or sensitive data in your workspace, and create a backup of your files outside of IBM Cloud Shell. You're responsible for your data, and your backup can help you recover in case an outage or data loss occurs for any reason. For more information about how to transfer files to create a backup, see Transferring multiple files.

What to do if a region is not available

If a region is not available, Cloud Shell indicates that the current region or location is not available. You must manually restore the files and data from a previous backup before you can access them.

To select an available region and manually restore your data and files, complete these steps:

1. In Cloud Shell, find the current location in the Cloud Shell menu bar, and click **Change**.
2. Select the location that you want to switch to, and click **OK**. For more information, see Changing locations.
3. Upload the previously backed-up files and data into the new region or location. For more information, see Uploading files.

Service dependency map for IBM Cloud Shell

If a service depends on other IBM Cloud services, there can be impacts if any of the dependent services are having issues. The dependency severity indicates the impact to the service when the dependency is down.

Critical

: When the the dependency is down, the service is down.

Significant

: When the dependency is down, the service features are impacted.

Medium

: When the dependency is down, the service might be impacted and a workaround is possible.

Minimal

: When the dependency is down, the main service features are not impacted.

The following table provides the dependency listing of this service following a standard deployment.

Dependencies	Dependency impacts	Customer provided	Control or data plane	Location of dependency
IBM Cloud Internet Services	Availability, Change management, Disaster recovery, Security compliance	No	Both	Same region
IBM Cloud Virtual Servers	Availability, Change management, Disaster recovery	No	Both	Same region
IBM Cloud Identity and Access Management	Availability, Change management, Security compliance	No	Both	Global
IBM Cloud Kubernetes Service and Red Hat OpenShift on IBM Cloud - containers-kubernetes	Availability, Change management, Disaster recovery	No	Both	Same region
IBM Cloud Bare Metal	Availability, Change management, Disaster recovery	No	Both	Same region
Akamai	Availability, Change management, Disaster recovery	No	Both	Global
IBM Cloud Object Storage	Availability, Change management, Disaster recovery, Security compliance	No	Both	Same region
IBM Cloud Databases - databases-for-etcd	Availability, Change management, Disaster recovery	No	Both	Same region
IBM Cloudant for IBM Cloud - cloudantnosqldb	Availability, Change management, Disaster recovery, Security compliance	No	Both	Same region
IBM Cloud Classic DNS Servers	Availability, Change management	No	Both	Same data center
IBM Cloud Classic Infrastructure Resource Management	Availability, Change management	No	Both	Global
IBM Cloud Public IP Address Management	Availability, Change management	No	Both	Same region
IBM Cloud Console	Availability, Operations	No	Both	Global

IBM Cloud Databases - databases-for-redis	Availability, Change management, Disaster recovery	No	Both	Same region

IBM Cloud Shell service dependency information - Critical dependencies

Dependencies	Dependency impacts	Customer provided	Control or data plane	Location of dependency
IBM Cloud Container Registry	Availability, Change management, Disaster recovery, Security compliance	No	Both	Same region
Synthetics	Availability, Operations	No	Both	Global
IBM Key Protect for IBM Cloud	Availability, Change management, Disaster recovery, Security compliance	No	Both	Same region
IBM Cloud Classic NTP Servers	Availability, Change management	No	Both	Same data center
IBM Cloud Activity Tracker event routing	Availability, Operations, Security compliance	No	Both	Same region
IBM Cloud Business Support Services	Availability	No	Both	Global
IBM Cloud Global Resource Catalog	Availability, Change management	No	Both	Global
Let's Encrypt	Availability, Change management, Security compliance	No	Both	Global
IBM Log Analysis Log Routing	Availability, Operations, Security compliance	No	Both	Same region
IBM Cloud Metrics Routing	Availability, Operations	No	Both	Same region
IBM Cloud Monitoring	Availability, Operations	No	Both	Same region
IBM Cloud Secrets Manager	Availability, Change management, Disaster recovery, Security compliance	No	Both	Same region
IBM Cloud Service Endpoints	Availability, Change management, Disaster recovery	No	Both	Same region
Segment	Availability	No	Both	Global

IBM Cloud Shell service dependency information - Significant dependencies

Dependencies	Dependency impacts	Customer provided	Control or data plane	Location of dependency
IBM Cloud CLI	Change management	No	Both	Global
TaaS Tekton Runners	Change management, Security compliance	No	Both	Global

IBM Cloud Shell service dependency information - Medium dependencies

Dependencies	Dependency impacts	Customer provided	Control or data plane	Location of dependency
Amplitude	Operations	No	Both	Global
SOS Inventory Management	Security compliance	No	Both	ibm-intranet

SOS Compliance Reporting	Security compliance	No	Both	ibm-intranet
SOS SIEM	Security compliance	No	Both	ibm-intranet
SOS Tenable	Security compliance	No	Both	ibm-intranet
SOS Health-Check, Vulnerability Scanning and Patching	Security compliance	No	Both	ibm-intranet
IBM Cloud Global Search and Tagging	Security compliance	No	Both	Global
IBM Log Analysis & IBM Cloud Activity Tracker	Operations, Security compliance	No	Both	Same region
OSS Platform	Operations	No	Both	Global

IBM Cloud Shell service dependency information - Minimal dependencies

This table can be used to answer the following questions:

- **What is the expected impact to the functions described?** Each severity tab in the table indicates the impact that your provisioned service might encounter if the dependency were to go offline. This means that the dependency high availability and disaster recovery influences the severity of the impact and therefore is used for general guidance to help you understand potential issues that might arise if the dependency was impacted by an incident.

  Note: Services that are regional are not impacted by a severe outage of a single availability zone because of the failover that is built in to default to another zone. For these occurrences, there might be a slight performance impact, if any, while the system fails over to the other location. This also applies to global services where the impact is lowered even more as it can fail over to other regions if necessary. This reduces the frequency at which these items might have the impact that is shown.

- **What services does my service depend on?** The Dependencies column lists the services. These are the major service to service dependencies including major internal dependencies that might not be visible externally.

- **What function does the dependency impact?** Functions include access management, availability, change management, configuration management, customer responsibility, disaster recovery, instance control, none, operations or security compliance. If the dependency goes offline, these functions might be impacted. Definitions for each available values are as follows:

  access management

  : Authentication, authorization and governance of the customer users access to the service and service instances.

  availability

  : Availability of the service and service instances.

  change management

  : Deployment, upgrade, patch, and so on of the service and service instances.

  configuration management

  : Deployment, upgrade, patch, and so on of the service and service instances.

  customer responsibility

  : Functions provides by customers to support specific service and service instances function. For example: IBM Key Protect for IBM Cloud instances provided by customer to support service BYOK encryption.

  disaster recovery

  : Backup, recovery, restart of the service and service instances in case of disruption.

  instance control

  : Creation, deletion, start, stop actions on lifecycle of the service instances.

  none

  : No function impacted.

operations

Monitoring, troubleshooting, etc of the service and service instances.

security compliance

Vulnerability management and other security and compliance management of the service and service instances.

- The **Customer provided** column will show if there is any dependency that has been provided by the customer to enable specific functionality. (for example: To properly configure and set up using BYOK into a service, the customer would provision a service like IBM Key Protect for IBM Cloud. But there may be other examples like this.) For details on how to enable the features and which services you need to provision, please see the documentation on the service.
- **Where do dependency services need to be deployed regarding my service?** In the Location of dependency column you can view if the dependency is located in the same region or deployed to a specific data center. You can use this data with the data in the Control or data plane column for a quick reference to identify if your data leaves the region or not in a standard setup.

  To find where your service can be deployed, see Service and infrastructure availability by location.

  > **Note:** The table shows a standard cloud deployment. If a special deployment is used like Fedramp or other region-bound deployment models, the data might differ from the details available in the table. Refer to the specific deployment that you are using for that information.

- **Where are the separate control plane and data plane located, if applicable?** Sometimes, the dependency might have a separate control plane and data plane. In these cases, there are separate lines that show the location in relation to the deployed customer instance of the service where these will be provisioned. The lines might have different impacts and different functions. See the Control or data plane column to understand what possible impact this type of outage might have.

  Same region means that the dependent services are in the same region as the provisioned instance. Other values might show data center or region names if the service must be used from a specific region, a specific availability zone, or set of availability zones. If a service is tied to a specific region or site, and the region goes offline, the service might go offline as well. It is recommended that you go through the high availability and disaster recovery documentation of the dependency to determine if there are any steps that you should take to mitigate these types of risks.

For more information about the policies that are related to the services, you can refer to the following resources:

- Service Level Agreement
- Shared responsibilities for using IBM Cloud products
- Service and infrastructure availability by location

Learning about IBM Cloud Shell architecture and workload isolation

Review the sample architecture for IBM® Cloud Shell, and learn more about different isolation levels. You can then choose the solution that best meets the requirements of the workloads that you want to run in the cloud.

IBM Cloud Shell architecture

IBM Cloud Shell is a public, multi-tenant regional service that is available in IBM Cloud®. With Cloud Shell, you can manage IBM Cloud resources and applications in a cloud-hosted shell environment from any web browser, with one click from the IBM Cloud console.

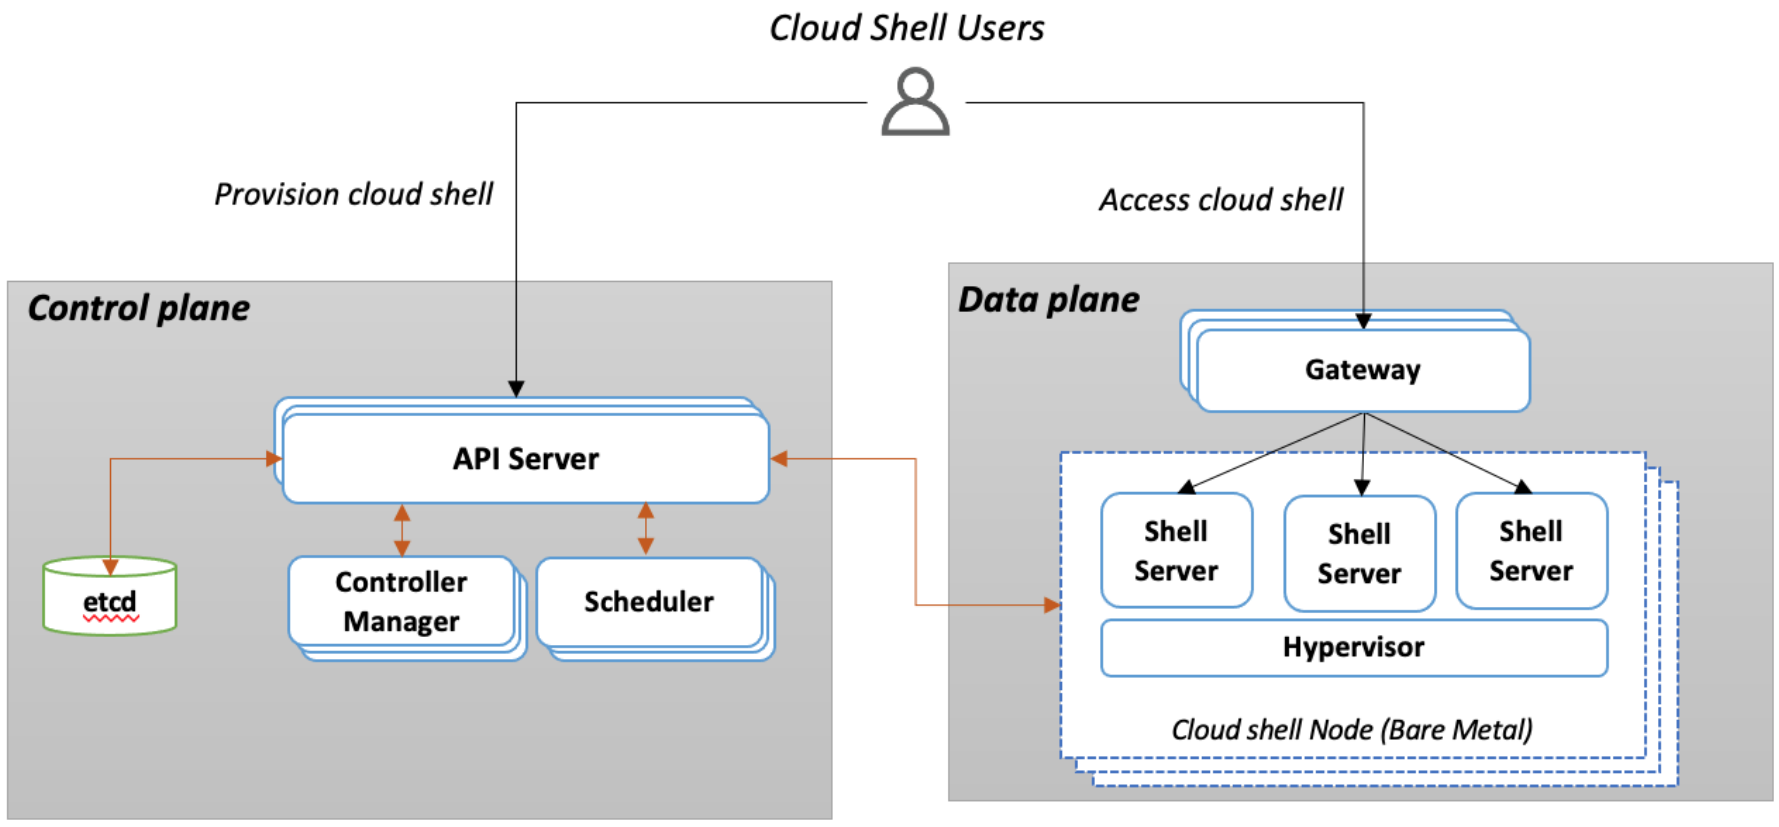

Cloud Shell sample architecture

The Cloud Shell control plane is responsible for provisioning Cloud Shell servers and managing the lifecycle.

- The API server provides an API interface to the Cloud Shell service.
- The controller manager manages the lifecycle of the Cloud Shell servers.
- The scheduler finds the best node that the Cloud Shell server is provisioned on.

The Cloud Shell data plane is hosting the user's Cloud Shell servers.

- The gateway is operating at the edge to manage the access and route the user's requests to the Cloud Shell server that is provisioned for the user.
- The Cloud Shell server is a virtual machine that is running on a bare metal server.

IBM Cloud Shell workload isolation

Each regional deployment of the IBM Cloud Shell serves multiple tenants. It is accessed through public endpoints. IBM keys encrypt all the data at rest. Data in transit is encrypted by using TLS.

The user's Cloud Shell server is running in a virtual machine. It is isolated from the other user's Cloud Shell server that is running on the same node. Cloud Shell servers that are running on the same node share physical resources such as CPU, memory, and I/O devices. The guest OS in an individual Cloud Shell server cannot detect any device other than the virtual devices that are made available to it.

The networks of the Cloud Shell servers are isolated from each other. Direct traffic is not allowed between Cloud Shell servers. The Cloud Shell server can access only the public internet.

IBM Cloud Shell session IP ranges

If you want to restrict access to IBM® Cloud Shell from your service, you can allowlist specific IP addresses or ranges of IP addresses on your deployment.

If you use an allowlist or blocklist to control connections in your environment, you can use the following IP information to allowlist access to Cloud Shell. Ensure that you allowlist all of the subnet ranges for the entire region where your deployments are.

You can configure the allowed and blocked IP list in your target service's IP list settings for traffic that is going to or coming from Cloud Shell.

You can use a Cloud Shell session as a terminal to use SSH to log in to your VSI. For example, if you use IBM Security Group to protect your VSI, you need to know the IP address range of Cloud Shell in a particular region.

Note: Cloud Shell does not have privileged access to IBM Cloud resources. It has the same access as the locally installed IBM Cloud CLI or the IBM Cloud console.

Current ingress IP ranges

Tables 1 - 3 list the IP ranges for each region.

Cloud Shell ingress IP ranges for Dallas (us-south)

Data center location	Subnet
Dallas 10	169.46.9.208/28 169.46.97.32/27 10.177.94.192/26
Dallas 12	169.48.197.32/28 169.48.231.32/27 10.185.217.192/26
Dallas 13	67.228.112.80/28 67.228.231.192/27 10.37.169.128/26

Cloud Shell ingress IP ranges for Dallas (us-south) caption-side=bottom"

Cloud Shell ingress IP ranges for eu-de (Frankfurt)

Data center location	Subnet
Frankfurt 2	169.50.22.224/27 10.134.69.192/26
Frankfurt 4	161.156.184.160/28 161.156.185.160/27 10.240.247.0/26
Frankfurt 5	149.81.65.224/28 149.81.112.128/27 10.123.216.64/26

Cloud Shell ingress IP ranges for eu-de (Frankfurt)

Installed plug-ins and tools

In addition to the full IBM Cloud® CLI, IBM® Cloud Shell includes preinstalled CLI plug-ins, packages, tools, and runtimes so that you can instantly start working on the command line.

For more information about the IBM Cloud Shell CLI plug-ins, see Extending IBM Cloud CLI with plug-ins.

IBM Cloud CLI plug-ins

IBM Cloud Shell includes all IBM Cloud CLI plug-ins, including but not limited to:

- IBM Cloud Activity Tracker CLI plug-in (`ibmcloud atracker`)
- IBM Analytics Engine CLI plug-in (`ibmcloud ae`)
- IBM Cloud App Configuration CLI plug-in (`ibmcloud ac`)
- IBM Cloud Catalogs management CLI plug-in (`ibmcloud catalog`)
- IBM Context-based restrictions CLI plug-in (`ibmcloud cbr`)
- IBM Cloud Databases CLI plug-in (`ibmcloud cdb`)
- IBM Cloudant CLI plug-in (`ibmcloud cloudant`)
- IBM Cloud Code Engine CLI plug-in (`ibmcloud ce`)
- IBM Code Risk Analyzer CLI plug-in (`ibmcloud cra`)
- IBM Cloud Container Registry CLI plug-in (`ibmcloud cr`)
- IBM Cloud DevOps Insights CLI plug-in (`ibmcloud doi`)
- IBM Direct Link CLI plug-in (`ibmcloud dl`)
- IBM Cloud Event Notifications CLI plug-in (`ibmcloud event-notifications`)
- IBM Cloud DNS Services CLI plug-in (`ibmcloud dns`)
- IBM Cloud Hyper Protect DBaaS CLI plug-in (`ibmcloud dbaas`)
- IBM Hyper Protect Crypto Services CLI plug-in (`ibmcloud tke`)
- IBM Hyper Protect Crypto Services Certificate manager CLI plug-in (`ibmcloud hpcs-cert-mgr`)
- IBM Hyper Protect Virtual Servers CLI plug-in (`ibmcloud hpvs`)
- IBM Cloud Internet Services CLI plug-in (`ibmcloud cis`)
- IBM Cloud Kubernetes Service CLI plug-in (`ibmcloud ks`)
- IBM Cloud Monitoring (ibmcloud monitoring) CLI plug-in (`ibmcloud monitoring`)
- IBM Cloud Object Storage CLI plug-in (`ibmcloud cos`)
- IBM Power Systems Virtual Server CLI plug-in (`ibmcloud pi`)
- IBM Cloud Privileged Access Gateway CLI plug-in (`ibmcloud gateway`)
- IBM Cloud Satellite CLI plug-in (`ibmcloud sat`)
- IBM Cloud Schematics CLI plug-in (`ibmcloud schematics`)
- IBM Cloud Secrets Manager CLI plug-in (`ibmcloud secrets-manager`)
- IBM Cloud Transit Gateway CLI plug-in (`ibmcloud tg`)
- IBM Cloud® Virtual Private Cloud CLI plug-in (`ibmcloud is`)
- IBM Event Streams CLI plug-in (`ibmcloud es`)
- IBM Key Protect CLI plug-in (`ibmcloud kp`)
- IBM Power® Virtual Server CLI plug-in (`ibmcloud pi`)
- IBM Cloud VPC CLI plug-in (`ibmcloud vpc-infrastructure`)
- IBM Watson Query CLI plug-in (`ibmcloud watson-query`)

Tip: The IBM Cloud CLI and plug-ins are regularly updated to the latest available versions. For the full list of installed plug-ins and their versions, run `ibmcloud plugin list` in Cloud Shell.

Tools

- Apache Ant (`ant`)
- Apache Maven (`maven`)
- Bazaar version control system (`bzr`)

- Calico CLI (`calicoctl`)
- cURL (`curl`)
- Docker (`docker`)
- etcd CLI (`etcdctl`)
- FIGlet text banner generator (`figlet`)
- Git CLI (`git`)
- GNU Automake (`automake`)
- GNU Compiler Collection (`gcc` , `gcov` , `gcov-tool`)
- GNU Wget (`wget`)
- Gradle CLI (`gradle`)
- Helm CLI (`helm`)
- Istio CLI (`istioctl`)
- jq JSON processor (`jq`)
- Knative CLI (`kn`)
- Kubernetes CLI (`kubectl`)
- Kubetail Kubernetes log aggregator (`kubetail`)
- Mercurial source content management tool (`hg`)
- Nano text editor (`nano`) and colors for `nano`
- nodemon Node.js development monitoring utility (`nodemon`)
- NPM Package Manager (`npm`)
- Node Version Manager (`nvm`)
- Operator SDK (`operator-sdk`)
- PostgeSQL client (`psql`)
- Red Hat® OpenShift CLI (`oc`)
- Red Hat OpenShift Do CLI (`odo`)
- Redis CLI (`redis-cli`)
- SoftLayer CLI (`slcli`)
- Terraform CLI (`terraform`) and IBM Cloud Provider plug-in for Terraform (`ibmcloud terraform`)
- tmux terminal multiplexer (`tmux`)
- UnZip extraction utility (`unzip`)
- Vim text editor (`vim`)
- XZ Utils (`xz` , `xzcat` , `xzcmp` , `xzdiff` , `xzegrep` , `xzfgrep` , `xzgrep` , `xzless` , `xzmore`)
- Yarn package manager (`yarn`)
- yq YAML processor (`yq`)
- Zip archive utility (`zip`)
- Simple Python Version Management (`pyenv`)

Packages and libraries

- `apt-transport-https`
- `bash-completion`
- `build-essential`
- `ca-certificates`
- `cron`
- `dnsutils`
- `gettext`
- `inetutils-ping`
- `libevent-dev`
- `libncurses5-dev`
- `locales`
- `musl`
- `net-tools`

- `pkg-config`
- `powerline-shell`
- Python `requests`
- `python-virtualenv`
- `python-setuptools`
- `python-pip`
- `python3-pip`
- `ruby-dev`
- `silversearcher-ag`
- `software-properties-common`
- `zlib1g-dev`

General IBM Cloud CLI (ibmcloud) commands

The IBM Cloud® Command Line Interface provides a set of commands that are grouped by namespace for users to interact with IBM Cloud.

The IBM Cloud CLI supports the following commands, including their names, arguments, options, prerequisites, descriptions, and examples.

Global prerequisites

The prerequisites for each command describe which actions are required before you run the command. The prerequisites can include one or more of the following actions:

Endpoint

: Use the `ibmcloud api` command to set an API endpoint.

Log in

: Use the `ibmcloud login` command to log in. If you log in with a federated ID, use the `--sso` option to authenticate with a one-time passcode. Or use the `--apikey` option to authenticate with an API key.

Target

: Use the `ibmcloud target` command to set or view the target account or region.

Global options

The following options are available for most commands in the IBM Cloud CLI. To check whether an option is available for a specific command, use the `-h, --help` option with the command.

`--output FORMAT`

Specifies an output format. Only JSON is supported.

Examples

Print available resource groups in JSON format:

```
ibmcloud resource groups --output json
```

`-q, --quiet`

Suppresses verbose messages. Prompt messages like `Getting information from... as ...` do not display if `-q, --quiet` is specified.

Examples

Print available resource groups in quiet mode:

```
ibmcloud resource groups -q
```

ibmcloud help

Displays the general help for first-level built-in commands and supported namespaces of IBM Cloud CLI, or the help for a specific built-in command or namespace.

```
$ ibmcloud help [COMMAND|NAMESPACE]
```

Prerequisites

None.

Command options

COMMAND|NAMESPACE

The command or namespace that help is displayed for. If not specified, the general help for IBM Cloud CLI is shown. Optional.

Examples

Display general help for the IBM Cloud CLI:

```
ibmcloud help
```

Display help for the `dev` command:

```
ibmcloud help dev
```

ibmcloud assist

Ask a question to the AI assistant that uses watsonx. For more information about the capabilities and limitations, see Getting help from the AI assistant.

```
$ ibmcloud assist "MY_QUERY"
```

Note: Content that is generated by the AI assistant might include mistakes or be incorrect.

Prerequisites

- Use the `ibmcloud login` command to log in. If you are logging in with a federated ID, use the `--sso` option to authenticate with a one-time passcode, or use the `--apikey` option to authenticate with an API key.

Command options

None.

Examples

```
ibmcloud assist "How do I update the CLI?"
```

Tip: If you're using quotations or exclamation marks as part of your query, make sure to escape each character by using `\`. For more information, see Using quotation marks with strings in IBM Cloud CLI.

ibmcloud version

Print the version of the IBM Cloud CLI.

```
$ ibmcloud version
```

Prerequisites

None.

Command options

None.

Examples

Print the version of the IBM Cloud CLI:

```
ibmcloud version
```

ibmcloud api

Set or view the IBM Cloud API endpoint.

```
$ ibmcloud api [API_ENDPOINT] [--unset] [--skip-ssl-validation] [--vpc]
```

Prerequisites

None.

Command options

API_ENDPOINT

The API endpoint that is targeted, for example, `https://cloud.ibm.com`. If both the `API_ENDPOINT` and `--unset` options aren't specified, the current API endpoint is displayed. Optional.

--skip-ssl-validation

Bypass SSL validation of HTTP requests. This option isn't recommended.

--vpc

Use a VPC connection for a private API endpoint. For more information about configuring to use a VPC connection for a private API endpoint, see Configuring a private endpoint gateway (required for VPC use).

--unset

Remove the API endpoint setting.

Examples

Set the API endpoint to `cloud.ibm.com`:

```
ibmcloud api cloud.ibm.com
```

Set the API endpoint to cloud.ibm.com and bypass SSL validation.

```
ibmcloud api https://cloud.ibm.com --skip-ssl-validation
```

View the current API endpoint:

```
ibmcloud api
```

Remove the API endpoint:

```
ibmcloud api --unset
```

ibmcloud config

Writes default values to the configuration file.

```
$ ibmcloud config --http-timeout TIMEOUT_IN_SECONDS | --trace (true|false|path/to/file) |
         --color (true|false) | --locale (LOCALE|CLEAR) | --check-version (true|false)
```

Prerequisites

None.

Command options

--check-version

Enable or disable CLI version checking. Valid values are `true` or `false`.

--color

Enable or disable color output. This option is disabled by default. Valid values are `true` or `false`.

--http-timeout

The timeout value for HTTP requests in seconds. The default value is 60 seconds.

--locale

Set a default locale. If no value is specified, the previous locale is deleted.

--sso-otp

Set the style of single sign-on (SSO) one-time passcode (OTP) retrieval. The default value is `manual` and requires user interaction. The `auto` setting will launch the default browser automatically and accept the token.

--trace

Trace HTTP requests to the terminal or specified file. Valid values are `true` or `false`.

Tip: You can specify only one of the options at a time.

Examples

Set the HTTP request timeout to 30 seconds:

```
ibmcloud config --http-timeout 30
```

Enable trace output for HTTP requests:

```
ibmcloud config --trace true
```

Trace HTTP requests to the `/home/usera/my_trace` file:

```
ibmcloud config --trace /home/usera/my_trace
```

Disable color output:

```
ibmcloud config --color false
```

Set the locale to `zh_Hans`:

```
ibmcloud config --locale zh_Hans
```

Clear the locale settings:

```
ibmcloud config --locale CLEAR
```

Enable automatic SSO one-time passcode acceptance:

```
ibmcloud config --sso-otp auto
```

ibmcloud login

Log in to the IBM Cloud CLI:

```
ibmcloud login [-a API_ENDPOINT] [--sso] [-u USERNAME] [-p PASSWORD] [--apikey KEY | @KEY_FILE] [--cr-token (TOKEN |
@CR_TOKEN_FILE) | --vpc-cri] [--profile PROFILE_ID | PROFILE_NAME | PROFILE_CRN] [-c (ACCOUNT_ID | ACCOUNT_OWNER_USER_ID) | --no-
account] [--accept] [-g (RESOURCE_GROUP_NAME | RESOURCE_GROUP_ID)] [-r REGION | --no-region] [--vpc]
```

Prerequisites

None.

Command options

-a API_ENDPOINT

The API endpoint. For example, `cloud.ibm.com`. Or, use `private.cloud.ibm.com` to log in by using a private endpoint. Using this flag saves the API endpoint to the configuration file.

--sso

Specify this option to log in with a federated ID. Using this option prompts you to authenticate with your single sign-on provider and enter a one-time passcode to log in.

-u USERNAME

The user name. Optional.

-p PASSWORD

The user password. Optional.

--apikey API_KEY or @API_KEY_FILE_PATH

The API key content or the path of an API key file that is indicated by the @ symbol.

--cr-token TOKEN or @CR_TOKEN_FILE_PATH

The compute resource token content or the path of a compute resource token file that is indicated by the @ symbol. If provided, the `--profile` flag, or `IBMCLOUD_CR_PROFILE` environment variable, must also be provided or set.

--vpc-cri

Log in as a VPC VSI compute resource identity. For more information, see Logging in as a Virtual Server Instance Compute Resource Identity.

--profile PROFILE_ID, PROFILE_NAME, or PROFILE_CRN

The name, ID, or CRN of the linked trusted IAM profile is used when you obtain the IAM access token. If provided, the `--cr-token` flag, `IBMCLOUD_CR_TOKEN` environment variable, or `--vpc-cri` flag must also be provided or set. If you're authenticating as a VPC VSI compute resource, specifying only a trusted profile CRN or ID is supported.

-c ACCOUNT_ID

The ID of the target account. This option is exclusive with the `--no-account` option.

--no-account

Forced login without the account. This option isn't recommended, and it is exclusive with the `-c` option.

--accept

Accept an invitation to join the targeted account. The account that is provided must be a valid account ID.

-g RESOURCE_GROUP

The name or ID of the target resource group. Optional.

-r REGION

The name of the target region. For example, `us-south` or `eu-gb`.

--no-region

Forced login without targeting a region.

--skip-ssl-validation

Bypass the SSL validation of HTTP requests. This option isn't recommended.

--vpc

Use a VPC connection for a private API endpoint. For more information about logging in by using a VPC connection for a private API endpoint, see Logging in to the CLI with a private endpoint. For information about configuring a VPC connection to use a private API endpoint, see Configuring a private endpoint gateway (required for VPC use).

Examples

Log in interactively:

```
ibmcloud login
```

Log in to a private endpoint:

```
ibmcloud login -a private.cloud.ibm.com
```

Two regions are supported: `us-south` and `us-east`.

Log in with a username and password, and set a target account:

```
ibmcloud login -u username -p password -c MyAccountID
```

Log in with federated ID, and set a target account:

```
ibmcloud login --sso -c MyAccountID
```

Use an API key with an associated account:

```
ibmcloud login --apikey api-key-string
```

```
ibmcloud login --apikey @filename
```

Use an API key with no associated account:

```
ibmcloud login --apikey api-key-string -c MyAccountID
```

```
ibmcloud login --apikey @fileName -c MyAccountID
```

Note: If the API key has an associated account, switching to another account isn't supported.

Log in as a specific user with a federated ID:

```
ibmcloud login --sso
```

Then, the CLI provides a URL link and prompts you for the passcode:

```
Get One Time Code from https://identity-1.us-south.iam.cloud.ibm.com/identity/passcode to proceed.
Open the URL in the default browser? [Y/n] >
```

Open the link in a browser to get a passcode. Enter the passcode in the console to log in.

In addition to using the URL link provided by the CLI, you can start the IBM Cloud console and log in. In the IBM Cloud console, to the **Avatar** icon > **Log in to CLI and API**. Copy the IBM Cloud CLI passcode command and enter it. Use this method to log in to your IBMid with your Google account or if you connected your cloud account with an IBM Cloud App ID instance.

Log in as an IBM Kubernetes Service Compute Resource linked to a trusted profile:

```
ibmcloud login --cr-token token-string --profile trusted_profile_name_id_or_crn
```

```
ibmcloud login --cr-token @filename --profile trusted_profile_name_id_or_crn
```

```
IBMCLOUD_CR_TOKEN=@filename ibmcloud login --profile trusted_profile_name_id_or_crn
```

```
IBMCLOUD_CR_TOKEN=@filename IBMCLOUD_CR_PROFILE=trusted_profile_name_id_or_crn ibmcloud login
```

For more information about logging in as an IBM Kubernetes Service compute resource, see Logging in with a Compute Resource token.

Log in as a VPC VSI Compute Resource by using the default trusted profile linked during instance provisioning:

```
ibmcloud login --vpc-cri
```

Log in as a VPC VSI Compute Resource linked to the specified trusted profile:

```
ibmcloud login --vpc-cri --profile trusted_profile_id_or_crn
```

```
IBMCLOUD_CR_PROFILE=trusted_profile_id_or_crn ibmcloud login --vpc-cri
```

For more information about logging in as a VPC VSI compute resource, see Logging in as a Virtual Server Instance Compute Resource Identity .

Accept an invitation to join a new account:

```
ibmcloud login -c TargetedAccountID --accept
```

ibmcloud logout

Log out of the CLI:

```
ibmcloud logout
```

Prerequisites

None.

ibmcloud regions

View the information for all regions on IBM Cloud:

```
ibmcloud regions
```

Prerequisites

Use the `ibmcloud api` command to set an API endpoint.

ibmcloud target

Set or view the target account or region:

```
$ ibmcloud target [-r REGION_NAME | --unset-region] [-c ACCOUNT_ID] [-g RESOURCE_GROUP | --unset-resource-group]
```

Prerequisites

- Use the `ibmcloud api` command to set an API endpoint.
- Use the `ibmcloud login` command to log in. If you are logging in with a federated ID, use the `--sso` option to authenticate with a one-time passcode, or use the `--apikey` option to authenticate with an API key.

Command options

-c ACCOUNT_ID

The ID of the target account. Optional.

-r REGION

The name of the target region, for example, us-south or eu-gb. Optional.

-g RESOURCE_GROUP

The name of the target resource group. Optional.

--unset-region

Clear the targeted region.

--unset-resource-group

Clear the targeted resource group.

Note: If none of the options are specified, the current account and region are displayed.

Examples

Set the current account:

```
ibmcloud target -c MyAccountID
```

Switch to a new region:

```
ibmcloud target -r eu-gb
```

View the current account and region:

```
ibmcloud target
```

ibmcloud update

Update the CLI to the most recent version:

```
$ ibmcloud update [-f]
```

Prerequisites

None.

Command options

-f

Force an update without confirmation. Root privilege is required.

General classic infrastructure service commands

Use classic infrastructure commands in the IBM Cloud CLI to configure and manage infrastructure services.

Run the `ibmcloud sl` command to see the list of available commands:

```
USAGE:
   ibmcloud sl command [arguments...] [options...]

COMMANDS:
   block           Classic infrastructure Block Storage
   call-api        Call arbitrary API endpoints
   file            Classic infrastructure File Storage
   dns             Classic infrastructure Domain Name System
   globalip        Classic infrastructure Global IP addresses
   hardware        Classic infrastructure hardware servers
   image           Classic infrastructure Compute images
   ipsec           Classic infrastructure IPSEC VPN
   order           Classic infrastructure Orders
   placement-group Classic infrastructure Placement Group
   security        Classic infrastructure SSH Keys and SSL Certificates
   securitygroup   Classic infrastructure network security groups
   subnet          Classic infrastructure Network subnets
   ticket          Classic infrastructure Manage Tickets
```

```
  user            Classic infrastructure Manage Users
  vlan            Classic infrastructure Network VLANs
  vs              Classic infrastructure Virtual Servers
  help            Print command usage message
```

To view help information about a command, run the following command:

```
$ ibmcloud sl [command] -h
```

For detailed information about each command, see the related reference topics in this documentation.

Note: The `ibmcloud sl init` command is no longer available as of CLI version `0.14`. To install the most recent CLI version, see Installing the stand-alone IBM Cloud CLI.

ibmcloud sl help

View help information for all commands to operate the classic infrastructure environment:

```
ibmcloud sl help
```

Managing Cloud Shell settings for account users

IBM® Cloud Shell settings are managed in the IBM Cloud® console. As an account owner or Cloud Shell administrator, you can control whether users in an account can access Cloud Shell, and you can select the location availability for an account.

IBM Cloud Shell is a cloud-based shell workspace that you can access through your browser. Cloud Shell is preconfigured with the full IBM Cloud CLI, plug-ins, and tools that you can use to manage apps, resources, and infrastructure. For more information, see Getting started with IBM Cloud Shell.

Before you begin

Only account owners, users assigned the Administrator role for the Cloud Shell account management service, or users assigned the Administrator role on all account management services can change the Cloud Shell settings. To assign this access to a user in your account, complete the following steps:

1. In the IBM Cloud console, go to **Manage** > **Access (IAM)**, and select **Users**.
2. On the Users page, select the user that you want to assign the role to.
3. On the individual user's page, click the **Access** tab, and then click **Assign access**.
4. Select the service **IBM Cloud Shell**.
5. For the role, select **Administrator**, and then click **Review**. For more information, see IAM roles.
6. Click **Add** to add your policy configuration to your policy summary.
7. Click **Assign**.

For more information, see the IAM roles and actions for the IBM Cloud Shell account management service.

Enabling or disabling Cloud Shell for an account

By default, Cloud Shell is enabled for an account. As an account owner or user with the correct access, you can enable or disable Cloud Shell for users in the account.

When the Cloud Shell availability setting is enabled, Cloud Shell is available to all users in the account. If the setting is disabled, no users in the account can access Cloud Shell. The **IBM Cloud Shell** icon is disabled in the IBM Cloud console.

To enable or disable Cloud Shell for the account, complete the following steps:

1. In the IBM Cloud console, go to **Manage** > **Account**, and select **IBM Cloud Shell settings**.
2. Select the **Enabled** or **Disabled** toggle, and then click **Save changes**.

Enabling or disabling Cloud Shell locations for an account

By default, all locations for the account are enabled, and the nearest available location is selected. Users are routed to the nearest available location, such as Dallas (us-south) or Frankfurt (eu-de).

As an account owner or user with the correct access, you can select whether Cloud Shell is enabled only in specific locations for the account. To select Cloud Shell locations for the account, complete the following steps:

1. In the IBM Cloud console, go to **Manage** > **Account**, and select **IBM Cloud Shell settings**.
2. Ensure that Cloud Shell Availability is enabled.
3. Select the toggle for each location that you want to enable or disable for the account.
4. Optional: Select **Enable new locations by default** to automatically enable new locations when they are available. If this option is not selected, you must select the toggle for each new location that you want to enable as it becomes available.
5. Click **Save changes**.

Enabling or disabling Cloud Shell features for an account

Account owners or users with Cloud Shell administrator access can enable or disable Cloud Shell features for an account. By default, all features for the account are enabled. The feature settings apply only to the enabled Cloud Shell locations.

To enable or disable Cloud Shell features for the account, complete the following steps:

1. In the IBM Cloud console, go to **Manage** > **Account**, and select **IBM Cloud Shell settings**.
2. Select the toggle for the feature that you want to enable or disable for the account. For example, **File upload and download** and **Web preview**.
3. Optional: Select **Enable new features by default** to automatically allow new features to be enabled for the account when they are available. If this option is not selected, you must select the toggle for each new feature that you want to enable as it becomes available.

4. Click **Save changes**.

Assigning access to Cloud Shell and its features at a user level

An account administrator can grant specific users access to Cloud Shell and its features, such as **File upload and download** and **Web preview**, even if Cloud Shell settings are disabled at the account level.

The IAM policy can be applied to specific locations with different roles. The roles are used to control the access of specific Cloud Shell features.

The IAM policy takes priority and is active only if the Cloud Shell account setting is disabled. If the Cloud Shell account setting is enabled and the IAM policy is set, the IAM policy has no effect. In that scenario, all users in account can access Cloud Shell.

To assign Cloud Shell access to a particular user, complete the following steps:

1. In the IBM Cloud console, go to **Manage** > **Access (IAM)**, and select **Users**.
2. On the Users page, select the user that you want to assign the role to.
3. On the individual user's page, click the **Access polcies** tab, and then click **Assign access**.
4. For the service, select **IBM Cloud Shell**. Then, click **Next**.
5. Scope the access to **Specific resources**. Select a location to enable the features in. Then, click **Next**.
6. Select one or more roles to assign to the user. For example, if you want to enable the **File Upload** and **File Download** features for the user, select the **File Manager** role. For more information, see IAM roles.
7. Click **Review**.
8. Click **Add** to add your policy configuration to your policy summary.
9. Click **Assign**.

When the user logs in to their IBM Cloud account, the user now has access to Cloud Shell and the file management features within Cloud Shell.

FAQs for IBM Cloud Shell

FAQs for IBM® Cloud Shell include questions about Cloud Shell sessions and its included tools and plug-ins. To find all FAQs for IBM Cloud®, see the FAQ library.

What type of shell does Cloud Shell use?

IBM Cloud Shell uses a Red Hat® Linux® bash shell.

What tools and plug-ins are included in Cloud Shell?

IBM Cloud Shell includes all available IBM Cloud CLI plug-ins and dozens of tools, packages, and runtimes. For the full list, see Installed plug-ins and tools.

What browsers does Cloud Shell support?

To work in Cloud Shell, you need to use one of the IBM Cloud supported browsers. For more information, see What are the IBM Cloud prerequisites? If you use a browser that is not supported, you might see blank screens or other display problems when you use Cloud Shell.

How do I copy and paste text in Cloud Shell?

To copy text in Cloud Shell, select the text that you want to copy, and then do one of the following actions:

- In your browser menu, click **Edit**, and then click **Copy** or **Paste**.
- For macOS, use the keyboard shortcuts Cmd+C and Cmd+V.
- For Linux®, use Ctrl+Insert for copy and Shift+Insert for paste.
- For Windows®, use Ctrl+Insert for copy and either Shift+Insert or Ctrl+Shift+V for paste.
- Select text and then right-click to use the menu, and select **Copy** or **Paste**.

Firefox and Internet Explorer might not support clipboard permissions properly.

How long can I use Cloud Shell?

You can use Cloud Shell for up to 50 hours within a single week. If you reach this limit, all Cloud Shell sessions are closed, and any data in your workspace is deleted. But don't worry, you can still access Cloud Shell only in 5-minute increments until the week resets.

How can I audit Cloud Shell usage?

Cloud Shell generates IBM Cloud Activity Tracker events for your sessions and commands. You can analyze these events by using the Activity Tracker service. For more information, see Activity Tracker events for IBM Cloud Shell.

How can I delete my data in Cloud Shell?

Your data in Cloud Shell is automatically deleted when Cloud Shell is closed after inactivity or reaching the usage limits.

How can I edit files in Cloud Shell?

Cloud Shell includes two text editors, Vim (`vim`) and Nano (`nano`). You can use either of these editors to work with files in Cloud Shell.

How can I customize my Cloud Shell?

As with any bash shell, you can modify the `.bashrc` file in your home directory to run commands or scripts every time a session starts. For example, you might set command aliases or environment variables that you often use. Because your home directory space is temporary, you need to edit the `.bashrc` file each time Cloud Shell restarts.

Be careful when you edit these values because you can introduce errors that cause your sessions to not start. Don't change the `CLOUDSHELL` , `BLUEMIX_HOME` , `ACCOUNT_ID` , and `SESSION_NAME` environment variables, because they're required for your Cloud Shell environment to work.

How can I switch accounts in Cloud Shell?

To switch the default account for all sessions, close Cloud Shell, switch the account in the IBM Cloud console menu bar, and then reopen Cloud Shell.

Can I use `sudo` to run commands?

Cloud Shell is a restricted shell, so `sudo` isn't supported in Cloud Shell.

How do I manage access and locations for Cloud Shell?

Account owners or Cloud Shell administrators can control whether users in an account can access Cloud Shell, and they can select the location for storing user session data. For more information, see the following topics:

- Updating Cloud Shell settings
- Actions and roles for account management services

Troubleshooting for IBM Cloud Shell

Use these tips to help troubleshoot problems that you might run into when you use IBM® Cloud Shell.

Tip: Running into issues that you can't fix? We want to hear about it. Submit feedback about Cloud Shell through our UI. Or come chat with the development team on IBM Cloud Dev Tools Slack. After you request your invitation, sign in and join the `#ask-your-question` channel.

Why can't my sessions start?

What's happening

When you try to start Cloud Shell or restart it after you were idle or your connection was lost, the following message is displayed:

`Your session couldn't be started.` or `Session could not be started. Please wait for a few minutes, then restart the cloud shell and try again.`

How to fix it

Refresh your Cloud Shell browser tab or window. If Cloud Shell still can't start, try the following possible solutions:

- Restart from the Cloud Shell menu.
- Check whether the IP address restriction is defined for the account or the user. If the IP address restriction is defined, add the Cloud Shell ingress IP ranges to log in from the Cloud Shell service. If the IP address restriction is not defined and you continue to see this issue, create a support case.
- Log out of IBM Cloud®. Clear your browser cache and cookies to remove your preferences and then log in again and open Cloud Shell.
- Check whether Concurrent sessions are set. If so, ensure that the number of logged in sessions for the account does not exceed the **Limit sessions** value.

Why can't I work with my Kubernetes clusters from my session?

What's happening

You want to work with your IBM Cloud® Kubernetes Service clusters, but when you run a command such as `kubectl get pods`, the following error is displayed:

```
$ $ kubectl get pods
The connection to the server localhost:8080 was refused - did you specify the correct host or port?
```

Why it's happening

The cluster isn't currently set as the context. As with your local development environment, the cluster context must be set for each individual session.

How to fix it

Set the cluster as the context in your session as described in Installing the CLI.

What do I do if I changed my `.bashrc` file and my sessions don't work?

What's happening

You customized your Cloud Shell sessions by editing the `.bashrc` file, and now your sessions don't open. As a result, you can't work in Cloud Shell.

Why it's happening

Some code in your `.bashrc` file isn't working correctly, and it's interfering with your sessions' ability to initialize.

How to fix it

If you're able to run commands from an existing open session, download any files that you want to keep. Then, restart Cloud Shell by going to the Cloud Shell menu and clicking **Restart**.

Why do I keep losing my connection to Cloud Shell?

What’s happening

You're working in Cloud Shell when your session closes with the following message:

```
The connection to your session was lost.
```

Why it’s happening

You might see this error for a few reasons:

- Your internet connection was unstable.
- You reached the maximum capacity for your Cloud Shell temporary storage.
- Rarely, the Cloud Shell region you were working in is unavailable.

How to fix it

If your connection was lost because of a temporary hiccup in your internet connection, refresh the browser window to reconnect. If you continue to be disconnected, restart Cloud Shell. Restarting IBM Cloud Shell removes all files in your workspace, so be sure to download any files that you want to keep.

If a region is unavailable, restarting Cloud Shell automatically loads it in an available region if you didn't change the Cloud Shell location from the default. If you did manually change to a different location, clear your browser cache and cookies to remove your preferences, and then open Cloud Shell.

Why does Cloud Shell hang, close, or restart?

What’s happening

You're working in Cloud Shell when your session hangs or closes with no message.

Why it’s happening

Cloud Shell might hang, experience slow performance, or close with intensive resource usage. If too many resources are used, the Cloud Shell environment might restart. When Cloud Shell closes, any data in your workspace is deleted.

Known issues and limitations

IBM® Cloud Shell includes the following known issues and limitations that might impact your experience.

Known issues

- Your Cloud Shell workspace includes 500 MB of temporary storage. If you use all of your storage space, your connection to Cloud Shell is lost. To access Cloud Shell, you must restart it, which removes all of your files.
- Adding files outside of your `/home/<user-name>` home directory can cause Cloud Shell to close, particularly if the files are large.
- You can mount (or mount by using a volume switch) your home directory or subdirectories of your home directory to your Docker container. If you mount to another location, an error results.

Limitations

- IBM Cloud Shell runs in the Dallas (`us-south`) and Frankfurt (`eu-de`) regions only. You can access resources in any region.
- Usage is limited to 50 hours of weekly use. For more information, see Tracking your usage.
- IBM Cloud Shell does not support trusted profiles. If you start IBM Cloud Shell and see an `Internal Error` , ensure that you're not using the IAM Trusted profiles feature to log in to your IBM Cloud account.

Release notes for IBM Cloud Shell

Use the release notes to learn about the latest changes to IBM® Cloud Shell.

Tip: The IBM Cloud CLI and plug-ins are regularly updated to the most recent available versions. These updates are not usually detailed in these release notes.

28 March 2024

Tokyo region is deprecated for Cloud Shell

The Tokyo region for Cloud Shell is deprecated and will no longer be supported by Cloud Shell as of 2 July 2024. Cloud Shell users who selected Tokyo as their region will automatically be redirected to the next geographically closest location (Dallas or Frankfurt). If you prefer, you can switch your location yourself within Cloud Shell. For more information, see Changing locations.

If an account's Cloud Shell settings have enabled only Tokyo Cloud Shell availability, then the account owner or user with the correct access must change their Cloud Shell settings to another listed location before 2 July 2024 to prevent any disruption in Cloud Shell performance. For more information, see Enabling or disabling Cloud Shell locations for an account .

17 May 2023

Preinstalled pyenv is available

Simple Python Version Management (`pyenv`) is now available on Cloud Shell. A preinstalled Python 3.11.2 is available by using `pyenv`.

31 March 2023

ibmcloud dev commands not available on Docker

When you deploy or run your application on Docker, the `ibmcloud dev` app commands are no longer available, including `ibmcloud dev build` and `ibmcloud dev run`. However, you can still use Docker commands in Cloud Shell.

21 September 2022

Docker in Docker is supported

Cloud Shell now supports Docker in Docker, which includes the ability to build and run with `ibmcloud dev build` and `ibmcloud dev run` when deploying or running your application on Docker.

3 August 2022

Redis 6 is supported

The supported version of Redis is now version 6.

14 April 2022

Helm version 3 is now the default version

Helm version 3 is now the default version. When you run `helm` on Cloud Shell, Helm version 3 is the version that runs. To run Helm version 2, type `helm2`. The `helm3` alias is also available if you use Helm version 3 in your scripts.

18 December 2020

Features can be enabled or disabled for an account

Account owners or users with Cloud Shell administrator access can enable or disable Cloud Shell features for an account. The available features in this release are **File upload and download** and **Web preview**. The feature settings apply only to the enabled Cloud Shell locations. For more information, see Enabling or disabling Cloud Shell features for an account.

Specific users can be granted access to features

An account administrator can grant specific users access to Cloud Shell and its features, even if Cloud Shell settings are disabled at the account level. For more information, see Assigning access to Cloud Shell and its features at a user level.

New service roles are available

The following new service roles are available:

- Cloud Operator
- Cloud Developer
- File Manager

For more information, see IAM roles and actions.

Red Hat Linux bash shell is now used

IBM Cloud Shell now uses a Red Hat™ Linux™ bash shell instead of a x86-64 Ubuntu Linux™ bash shell.

18 September 2020

Cloud Shell settings are managed from the IBM Cloud console

Account owners or users with Cloud Shell administrator access can manage Cloud Shell settings from the IBM Cloud console. For more information, see the following topics:

- Updating Cloud Shell settings
- Actions and roles for account management services

New and updated tools are available

The following tools are new with this Cloud Shell update:

- Terraform version switch (`tfswitch`) version 0.8.832
- Tekton CLI (`tkn`) version 0.12.0

The following tools are updated with this Cloud Shell update:

- `kn` upgrade to 0.14.0
- The IBM Cloud CLI (`ibmcloud`) is upgraded to version 1.2.1, and the CLI plug-ins are upgraded to the most recent versions

12 August 2020

Weekly usage quote increased to 50 hours

Increased the weekly usage quota from 30 hours to 50 hours so you have more time to use Cloud Shell each week. For more information about the usage quota, see Tracking your usage.

Separate Cloud Shell workspaces are available for each account

Changed Cloud Shell to have separate workspaces for each account. Previously, workspaces in each region were set at the user level, which meant that your workspace was shared across all of your accounts. For more information about workspaces, see Your Cloud Shell workspace.

Cookie preferences moved to the Privacy option

Moved the cookie preferences from the in-page button to the **Privacy** option in the Cloud Shell menu.

Cloud Shell server image is version 1.0.1

Updated the Cloud Shell server image to version 1.0.1, which includes the following changes:

- Added the Operator SDK (`operator-sdk`) for managing Kubernetes applications at version 0.19.0.
- Added the Mercurial source content management tool (`hg`) at version 4.4.1.
- Added the Bazaar version control system and CLI (`bzr`) at version 2.7.0.
- Updated the Python runtime from version 3.5 to 3.6.
- Updated the IBM Cloud Provider plug-in for Terraform to 1.8.1

24 June 2020

IBM Cloud Shell is now generally available (GA)

With this release, Cloud Shell includes the following updates:

- Increased service level availability to 99.99% and hardened disaster recovery plan to meet IBM Cloud requirements. For more information, see Understanding high availability and disaster recovery for Cloud Shell.
- Added support for inputting and viewing double-byte characters on the command line in your sessions. This enables you to work in national languages such as Japanese, Simplified Chinese, or Traditional Chinese.
- Added support for the Tokyo (`jp-tok`) region.
- Added cookie preferences so you can choose which cookies Cloud Shell saves. Cloud Shell relies on personalization cookies to remember your region preferences.
- Added a link to these release notes from the About IBM Cloud Shell window, so you can find them even more easily in the future.
- Updated the Cloud Shell server image to version 0.7.1, which includes the following changes:
  - Updated the ISTIO CLI (`istioctl`) to 1.5.4.
  - Added auto-complete for Red Hat OpenShift CLI (`oc`).

15 May 2020

Server image is version 0.6.6

Updated the Cloud Shell server image to version 0.6.6:

- Updated Red Hat® OpenShift CLI (`oc`) to 4.3.
- Updated the IBM Cloud CLI and all plug-ins to the latest versions.
- Removed Netcat networking utility (`netcat`).

18 April 2020

Preview web apps

Added the ability to preview web apps that you access through Cloud Shell. For more information, see Previewing apps.

Frankfurt (`eu-de`) region is supported

Added support for running Cloud Shell in the Frankfurt (`eu-de`) region. Each region runs independently, so your workspace, sessions, and usage are separate. For more information, see Changing locations.

Time limit updates

Extended the time after which inactive sessions are closed from 30 minutes to an hour.

Extended the time after which workspace data is removed due to inactivity from 30 minutes to an hour after there are no open sessions.

Removed the 4-hour continuous usage limit. You can now use Cloud Shell for any length of time up to the weekly usage quota.

Cloud Shell server image is 0.6.3

Updated the Cloud Shell server image to version 0.6.3, which includes the following changes:

- Added GNU Automake (`automake`).
- Added GNU Compiler Collection (`gcc` , `gcov` , and `gcov-tool`).
- Downgraded the ISTIO CLI (`istioctl`) from 1.5 to 1.4.
- Updated the IBM Cloud CLI and all plug-ins to the most recent versions.

16 February 2020

Cloud Shell server image is 0.4.37

Updated the Cloud Shell server image to version 0.4.37. Changes to this image affect the tools and other capabilities that are available in each session.

Enabled auto-completion for the IBM Cloud CLI.

Added the community `kn` Knative CLI and removed the third-party `knctl` Knative CLI.

Added the IBM Watson CLI plug-in (`ibmcloud watson`).

Updated IBM Cloud CLI plug-ins to the most recent versions.

Updated the ISTIO CLI (`istoctl`) to version 1.4.4.

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America
2024-12-13