



dbt Cloud Architect Certification Exam Study Guide

How to use this study guide

This is the official study guide for the **dbt Cloud Architect Certification Exam** from the team at dbt Labs. While the guide suggests a sequence of courses and reading material, we recommend using it to supplement (rather than substitute) real-world use and experience with dbt.

The [exam overview](#) will provide high-level details on the format of the exam. We recommend being mindful of the number of questions and time constraints.

The [topic outline](#) will provide a clear list of the topics that are assessed on the exam. dbt subject matter experts (SMEs) used this topic outline to write and review all of the exam items that you will find on the exam.

The [sample exam questions](#) provide an example of the format to expect for the exam. The types of questions you can expect on the exam include

- Multiple-choice
- Fill-in-the-blank
- Matching
- [Hotspot](#)
- Build list
- [Discrete Option Multiple Choice \(DOMC\)](#)

The [learning path](#) will walk you through a suggested series of courses, readings, and documentation. We will also provide some guidance for the types of experience to build while using dbt on a real-world project.

Finally, the [additional resources](#) section will provide additional places to continue your learning.

Exam Overview

The [dbt Cloud Architect Certification Exam](#) assesses your ability to design secure, scalable dbt Cloud implementations, with a focus on environment orchestration, role-based access control, integrations with other tools, and collaborative development workflows aligned with best practices.

We recommend that you have SQL proficiency and at least six months of experience administrating an Enterprise dbt Cloud account.

Logistics

- **Duration:** 2 Hours
- **Format & Registration:** [online proctored](#)
- **Length:** 65 questions
- **Passing Score:** 65% or higher. You will know your score immediately after completion of the exam.
- **Price:** \$200*
- **Language:** Only English at this time
- **Certification Expiration:** The certification expires 2 years after the date awarded.
- **Supported browsers:** [Caveon Web browsers](#)

*The price applies only to the online proctored exam; there is no additional fee for Coalesce premium attendees. Discounts are available for dbt Labs SI Partners.

Scoring

The exam is scored on a point basis, with 1 point for each correct answer, and 0 points for incorrect answers. All questions are weighted equally.

An undisclosed number of unscored questions will be included on each exam. These are unmarked and indistinguishable from scored questions. Answers to these questions will be used for research purposes only, and will not count towards the score.

Retakes & Cancellations

If you do not pass the exam, you may schedule a retake. You will need to pay a registration fee for each retake. You can reschedule or cancel without penalty on Talview up to your scheduled exam. We will not issue refunds for no-shows.

Accommodations

Please contact [Talview](#) with any accommodation requests.

Topic Outline

The **dbt Cloud Architect Certification Exam** has been designed to assess the following topics and sub-topics.

Topic 1: Configuring dbt Cloud data warehouse connections

- Understanding how to connect the warehouse
- Configuring IP whitelist
- Creating and testing a connection for the project
- Authenticating through OAuth to access the data in IDE
- Adding Client ID and Secret for OAuth

Topic 2: Configuring dbt Cloud git connections

- Connecting the git repo to dbt
- Setting up integrations with git providers

Topic 3: Creating and maintaining dbt Cloud environments

- Understanding access control to different environments
- Determining when to use a service account
- Rotating key pair authentication via the API
- Understanding environment variables
- Creating new dbt Cloud deployment environment
- Setting default schema/dataset for environment
- Understanding custom branches and which to configure to environments
- Configuring dbt to allow deferral to other environments

Topic 4: Creating and maintaining job definitions

- Setup a CI job with deferral
- Understanding steps within a dbt job
- Scheduling a job to run on schedule
- Implementing run commands in the correct order
- Creating new dbt Cloud job
- Configuring optional settings such as environment variable overrides, threads, deferral, target name, and dbt version override

- Generating documentation on a job that populates the project's doc site
- Configuring jobs to be triggered after other dbt jobs (job chaining)
- Configuring Advanced CI
- Configuring self deferral
- Understanding when to use which type of job deferral

Topic 5: Configuring dbt Cloud security and licenses

- Creating Service tokens for API access
- Assigning permission sets
- Creating license mappings
- Adding and removing users
- Adding SSO application for dbt Cloud enterprise
- Creating and assigning RBAC

Topic 6: Setting up monitoring and alerting for jobs

- Setting up email notifications
- Using Webhooks for event-driven integrations with other systems

Topic 7: Setting up a dbt mesh and leveraging cross-project references

- Setting up additional dbt projects
- Understanding how environment types relate to cross-project references
- Utilizing model governance

Topic 8: Configuring and using dbt Explorer

- Using dbt Explorer to understand the current lineage, troubleshoot issues and optimise cost and performance
- Using dbt Explorer to find public models and cross-project references

Sample exam questions

Sample Exam Question 1

Why should you avoid using SSO-based connections for staging and production deployment environments in dbt Cloud?

- A. They restrict the ability to manage environment variables required for production deployments.
- B. They cannot connect to cloud data warehouses due to incompatible authentication protocols.
- C. They rely on authentication sessions that can expire or require re-authentication, leading to unexpected job failures in automated deployments.
- D. They can only be used by users with the Account Admin role.

Explanation: dbt Cloud supports authenticating via OAuth in Development environments only, where users can authorize their Development credentials using Single Sign-On. Once a user has authorized dbt Cloud with their data platform, via their identity provider, the user will receive (behind the scenes) an Access Token which can then be used to open a data platform connection and execute queries in the dbt Cloud IDE.

The lifetime of the token is dictated by the data platform integration. When a user's token expires, the user will need to re-authorize with the data platform to continue development in dbt Cloud. For this reason, it is best practice to utilize a service account for deployment credentials.

A is incorrect, because the method of authentication does not restrict the usage of environment variables.

B is incorrect because it is technically possible to connect to your warehouse this way, it is just not advised, so dbt Cloud does not allow it.

D is incorrect because SSO OAuth can be used by any user with a Developer license, regardless of permissions granted.

[Link to docs](#)

Sample Exam Question 2

You have configured a snapshot's `target_schema` to be `prod_snapshot`.

You want to limit who can write to this schema.

How can you allow developers to read from the `prod_schema` in dev but allow read and write privileges for the production job?

- A. Assign the developer the Job Admin role in dbt Cloud.
- B. Use an `on-run-end` to grant read and write privileges in the snapshot config.
- C. Use a `post-hook` to grant read and write privileges in the snapshot config.
- D. Use a service account with read and write privileges in the deployment environment settings.

Explanation: This question boils down to the fact that different users or service accounts should have different permissions to operate on different objects in the database. Developers should have read access to `prod_schema`, but not write access to `prod_snapshot`. The production deployment credentials should have read and write to `prod_schema` and `prod_snapshot`. This is handled succinctly by utilizing different credentials for each environment, and granting the appropriate permissions within the data platform to each user / service account.

A is incorrect because assigning the dbt Cloud Job Admin role will allow the developer to create and edit jobs, runs, environment variables, and data warehouse configs, but that does not affect their personal access to objects in the data database.

For answers B and C, utilizing a `post-hook` or `on-run-end` hook won't help, because the production deployment credentials need write access before it can create or alter the table in the `prod_snapshot` schema. You could grant privileges afterwards to certain groups, but that doesn't solve the ask at hand.

[Link to docs](#)

Sample Exam Question 3

Inspect this environment variable configuration:

Key	Project Default	Production	Development
DBT_ENV_VALUE	not_set	prod	dev

Within a job there is an override to this variable with the value finance.

Which value will be used when running the dbt Cloud job if this job configuration is removed?

- A. the job-level environment variable is required and cannot be removed
- B. the default argument supplied to the `env_var` Jinja function in code
- C. the dbt project default `not_set`
- D. the value assigned to the environment specified in the job description

Explanation: The Project Default column is recommended when you want to supply a catch-all default or add a project-wide token or secret. Values set at the environmental level take priority over the project-level default value, and as seen in the above example, dbt Cloud will interpret the variable value differently in your **Production** vs. **Development** environment.

A job inherits the values set at the environment level, but a job-level override will take precedence for the environment in which it runs. If you remove the job-level override, it will default back to the environment variable, in this case, the value **prod**.

A is incorrect because the job-level environment variable is not required and can be removed.

B is incorrect because the Jinja function default argument is lower in precedence than both Project Default and the environment.

C is incorrect because Project Default is lower in the order of precedence than the environment itself.

[Link to docs](#)

Sample Exam Question 4

Congratulations! Your company has expanded to Europe.

As part of this expansion, it plans to set up another data warehouse in the EU while leveraging the same Enterprise dbt account across regions.

What should you do to achieve this?

- A. configure dbt Mesh Cross Platform to connect to the EU warehouse
- B. upgrade your account to Business Critical and establish a PrivateLink connection
- C. purchase another account for the EU region
- D. configure a connection to the EU — no further action is required

Explanation: In a multi-Region Setup, it would be prudent to set up two separate PrivateLink endpoints in the regions in which you are established. You will then configure multiple projects to use the respective private endpoints. This helps you maintain data sovereignty (EU data stays in EU) and isolation (no data traverses public networks).

A is incorrect because you do not want to set up Cross Platform mesh, as you will want data in EU to stay segregated from the other region's data.

B is correct because you will need PrivateLink to meet data residency/compliance goals (e.g., GDPR) and this is only possible with Business Critical Enterprise plan or higher.

C is incorrect because you will not need to purchase another account.

D is incorrect, because there is further design and action required to make this work.

[Link to docs](#)

Learning Path:

This is not the **only** way to prepare for the exam, but just one recommended path for an advanced dbt user to prepare for the exam. Each checkpoint provides a logical set of courses to complete, readings and documentation to digest, and real-world experience to seek out. We recommend this order, but things can be reorganized based on your learning preferences.

Checkpoint 0 - Prerequisites

dbt is a tool that brings together several different technical skills in one place. We recommend starting this path after you've obtained the badge from the [dbt Fundamentals course](#).

Additionally, being aware of the [Analytics Development Lifecycle \(ADLC\)](#) provides an excellent implementation framework for architecting your dbt Cloud Architecture.

Checkpoint 1 - Setting up dbt Cloud Connections

Resources:

- Courses:
 - dbt Cloud and Databricks for Admins
<https://learn.getdbt.com/learn/course/dbt-cloud-and-databricks-for-admins>
 - dbt Cloud and BigQuery for Admins
<https://learn.getdbt.com/learn/course/dbt-cloud-and-bigquery-for-admins>
 - dbt Cloud and Snowflake for Admins
<https://learn.getdbt.com/learn/course/dbt-cloud-and-snowflake-for-admins>
 - dbt Cloud and Redshift for Admins
<https://learn.getdbt.com/courses/dbt-cloud-and-redshift-for-admins>
- Readings:
 - dbt Cloud Architecture
<https://docs.getdbt.com/docs/cloud/about-cloud/architecture>
 - Maximum override: Configuring unique connections in dbt Cloud
<https://docs.getdbt.com/blog/configuring-unique-connections-in-dbt-cloud>
 - Version control basics
<https://docs.getdbt.com/docs/collaborate/git/version-control-basics>
 - Getting Started with git Branching Strategies and dbt
<https://docs.getdbt.com/blog/git-branching-strategies-with-dbt>

-
- Documentation:
 - About data platform connections
<https://docs.getdbt.com/docs/cloud/connect-data-platform/about-connections>
 - Supported data platforms
<https://docs.getdbt.com/docs/supported-data-platforms>
 - Access, Regions & IP addresses
<https://docs.getdbt.com/docs/cloud/about-cloud/access-regions-ip-addresses>
 - Set up Snowflake OAuth
<https://docs.getdbt.com/docs/cloud/manage-access/set-up-snowflake-oauth>
 - Set up Databricks OAuth
<https://docs.getdbt.com/docs/cloud/manage-access/set-up-databricks-oauth>
 - Set up BigQuery OAuth
<https://docs.getdbt.com/docs/cloud/manage-access/set-up-bigquery-oauth>
 - Set up external OAuth
<https://docs.getdbt.com/docs/cloud/manage-access/external-oauth>
 - Connect to GitHub
<https://docs.getdbt.com/docs/cloud/git/connect-github>
 - Connect to GitLab
<https://docs.getdbt.com/docs/cloud/git/connect-gitlab>
 - Connect to Azure DevOps
<https://docs.getdbt.com/docs/cloud/git/connect-azure-devops>
 - Connect with Git clone
<https://docs.getdbt.com/docs/cloud/git/import-a-project-by-git-url>
 - Experience
 - Understanding how to connect the warehouse including configuring IP whitelist
 - Authenticating through OAuth to access the data in IDE
 - Version control basics

Checkpoint 2 - Configuring and Managing Projects

Resources:

- Courses:
 - Advanced Deployment
<https://learn.getdbt.com/learn/course/advanced-deployment>
- Readings:
 - dbt Cloud environments
<https://docs.getdbt.com/docs/dbt-cloud-environments>
 - More time coding, less time waiting: Mastering defer in dbt
<https://docs.getdbt.com/blog/defer-to-prod>
 - To defer or to clone, that is the question
<https://docs.getdbt.com/blog/to-defer-or-to-clone>
- Documentation:
 - About user access
<https://docs.getdbt.com/docs/cloud/manage-access/about-user-access>
 - About environment-level permissions
<https://docs.getdbt.com/docs/cloud/manage-access/environment-permissions>
 - Environment Variables
<https://docs.getdbt.com/docs/build/environment-variables>
 - Custom schemas
<https://docs.getdbt.com/docs/build/custom-schemas>
 - Custom target names
<https://docs.getdbt.com/docs/build/custom-target-names>
 - Customize dbt models database, schema and alias
<https://docs.getdbt.com/guides/customize-schema-alias>
 - Threads
<https://docs.getdbt.com/docs/running-a-dbt-project/using-threads>
 - Deployment environments
<https://docs.getdbt.com/docs/deploy/deploy-environments>
 - How do I use the 'Custom Branch' settings in a dbt Cloud Environment?
<https://docs.getdbt.com/faqs/Environments/custom-branch-settings>
 - Jobs in dbt Cloud
<https://docs.getdbt.com/docs/deploy/jobs>
 - Customizing CI/CD with custom pipelines
<https://docs.getdbt.com/guides/custom-cicd-pipelines>



-
- Get started with Continuous Integration tests
<https://docs.getdbt.com/guides/set-up-ci>
 - Advanced CI
<https://docs.getdbt.com/docs/deploy/advanced-ci>
 - Using defer in dbt Cloud
<https://docs.getdbt.com/docs/cloud/about-cloud-develop-defer>
 - Experience
 - Understanding and configuring access control across different environments and projects
 - Determining when to use a service account and rotating key pair authentication via the API
 - Creating new dbt Cloud deployment environment with customisations using environment variables, custom branches, custom schemas/datasets and deferrals
 - Setting up different dbt jobs (scheduled, continuous integration, merge or via APIs)
 - Configuring Advanced CI

Checkpoint 3 - Security & Monitoring

Resources:

- Courses:
 - Webhooks
<https://learn.getdbt.com/courses/webhooks>
 - Role-based access control: dbt and Okta
<https://learn.getdbt.com/learn/video/role-based-access-control-okta>
 - SSO dbt Cloud and Entra ID
<https://learn.getdbt.com/learn/video/sso-dbt-cloud-and-entra-id>
 - SSO dbt Cloud and Google Workspace
<https://learn.getdbt.com/learn/video/sso-dbt-cloud-and-google-workspace>
 - SSO dbt Cloud and Okta
<https://learn.getdbt.com/learn/video/sso-dbt-cloud-and-okta>
- Readings:
 - Single sign-on (SSO) Overview
<https://docs.getdbt.com/docs/cloud/manage-access/sso-overview>

-
- Documentation:
 - Service account tokens
<https://docs.getdbt.com/docs/dbt-cloud-apis/service-tokens>
 - Users and licenses
<https://docs.getdbt.com/docs/cloud/manage-access/seats-and-users>
 - Self-service Team account permissions
<https://docs.getdbt.com/docs/cloud/manage-access/self-service-permissions>
 - Enterprise permissions
<https://docs.getdbt.com/docs/cloud/manage-access/enterprise-permissions>
 - About user access
<https://docs.getdbt.com/docs/cloud/manage-access/about-user-access>
 - Migrating to Auth0 for SSO
<https://docs.getdbt.com/docs/cloud/manage-access/auth0-migration>
 - Set up SSO with SAML 2.0
<https://docs.getdbt.com/docs/cloud/manage-access/set-up-sso-saml-2.0>
 - Set up SSO with Okta
<https://docs.getdbt.com/docs/cloud/manage-access/set-up-sso-okta>
 - Set up SSO with Google Workspace
<https://docs.getdbt.com/docs/cloud/manage-access/set-up-sso-google-workspace>
 - Set up SSO with Microsoft Entra ID
<https://docs.getdbt.com/docs/cloud/manage-access/set-up-sso-microsoft-entra-id>
 - Set up SCIM
<https://docs.getdbt.com/docs/cloud/manage-access/scim>
 - Job notifications
<https://docs.getdbt.com/docs/deploy/job-notifications>
 - Webhooks for your jobs
<https://docs.getdbt.com/docs/deploy/webhooks>
 - Experience
 - Creating Service tokens for API access
 - Assigning permission sets, creating license mappings and managing users
 - Adding SSO application for dbt Cloud enterprise, including creating and assigning role-based access control (RBAC)
 - Setting up email notifications
 - Using Webhooks for event-driven integrations with other systems

Checkpoint 4 - Governance and Discovery

Resources:

- Courses:
 - dbt Mesh
<https://learn.getdbt.com/learn/course/dbt-mesh>
 - dbt Explorer
<https://learn.getdbt.com/learn/course/dbt-explorer>
- Readings:
 - How Hybrid Mesh unlocks dbt collaboration at scale
<https://docs.getdbt.com/blog/hybrid-mesh>
 - Column-Level Lineage, Model Performance, and Recommendations: ship trusted data products with dbt Explorer
<https://docs.getdbt.com/blog/dbt-explorer>
 - Implementing your mesh plan
<https://docs.getdbt.com/best-practices/how-we-mesh/mesh-4-implementation>
- Documentation:
 - Intro to dbt Mesh
<https://docs.getdbt.com/best-practices/how-we-mesh/mesh-1-intro>
 - Project dependencies
<https://docs.getdbt.com/docs/collaborate/govern/project-dependencies#how-to-write-cross-project-ref>
 - Model access
<https://docs.getdbt.com/docs/collaborate/govern/model-access>
 - Discover data with dbt Explorer
<https://docs.getdbt.com/docs/collaborate/explore-projects>
- Experience
 - Setting up additional dbt projects with cross-project references
 - Using dbt Explorer to understand the current lineage, troubleshoot issues and optimise cost and performance
 - Using dbt Explorer to find public models and cross-project references

Additional Resources:

- dbt Slack
 - #dbt-certification
 - #learn-on-demand
 - #advice-dbt-for-power-users
 - #dbt-deployment-and-orchestration
- If you are a dbt Labs partner or enterprise client, contact your partner manager or account team for additional benefits
- Want to use email? Contact certification@dbtlabs.com