



## Top 20 ETL Interview Questions and Answers

Share on your Social Media



# Top 20 ETL Interview Questions and Answers

Published On: June 3, 2024

## Top 20 ETL Interview Questions and Answers

Are you getting ready for a job interview in data engineering or ETL? With this collection of the most common **ETL interview questions and answers** on a range of ETL tools, you may be ready for your upcoming ETL job interview.

[Download ETL Interview Questions PDF](#)

## ETL Interview Questions and Answers for Freshers

### 1. What is an ETL process?

ETL is abbreviated as extraction, transformation, and loading.

## Featured Articles

 **Want to know more about becoming an expert in IT?**

[Click Here to Get Started](#) >>

100% Placement Assurance

AUTHORISED CERTIFICATION PARTNER

IBM



Quick Enquiry

## Related Courses at SLA

- [ETL Training in OMR](#)
- [ETL Training in Chennai](#)

## Related Posts



### MEAN Stack Interview Questions and Answers

Published On: June 19, 2024

## 2. How many steps are involved in an ETL process?

In its most basic form, data extraction, transformation, and loading comprise the ETL process. Even so, the acronym represents a clear and sequential three-step process: extract, transform, and load.

## 3. What do initial and full loads mean?

The process of filling every data warehousing table for the first time is known as the first load in ETL. Depending on the volume of the data, all set records are loaded at once during full load, when the data is loaded for the first time. It would reload the most recent data and remove everything from the table.

## 4. What features do snapshots have, and what does their significance entail?

The read-only data that is kept in the master table is duplicated in snapshots.

To record changes to the master table, snapshots are stored on remote nodes and are frequently refreshed. They are duplicates of tables as well.

## 5. Which applications is PowerCenter compatible with?

SAP, Oracle Apps, PeopleSoft, and other ERP sources can be integrated with PowerCenter.

## 6. What distinguishes PowerMart from PowerCenter?

While Power Mart handles modest volumes of data, PowerCenter processes huge volumes of data.

## 7. Define views.

One or more tables' properties are used to build views. It is possible to update a view with a single table, but not one with many tables.

Introduction Since MEAN Stack combines several other applications as part of its functionality, it is...



### Top 15 Struts Interview Questions and Answers

Published On: June 18, 2024

Struts Interview Questions and Answers When it comes to developing Java web applications, Struts is...



### Top 20 C Sharp Interview Questions and Answers

Published On: June 17, 2024

C Sharp Interview Questions and Answers Microsoft created the general-purpose programming language C# together with...



### Top 20 VB.Net Interview Questions and Answers

Published On: June 17, 2024

VB.Net Interview Questions and Answers A wide range of applications, including

## 8. What does incremental load mean?

Applying dynamic modifications as needed within predetermined timeframes is known as incremental load.

## 9. What does a materialized view log entail, and what exactly is a materialized view, in your opinion?

A table that records modifications to the basis tables used in a materialized view is called a materialized view log. A pre-computed aggregate table with combined or condensed data from fact and dimension tables is called a materialized view.

## 10. What does partitioning in ETL mean?

In ETL, partitioning is the splitting of the transactions into smaller groups to boost efficiency.

## 11. Does PowerMart offer links to sources of ERP data?

Not at all! PowerMart does not provide ERP source connectors.

## 12. Describe the idea of data skewness in ETL procedures.

Data skewness is the term for an imbalance in the distribution of data among processing nodes or partitions, which can lead to issues with the performance of ETL processes. An imbalance occurs when certain keys or values occur far more frequently than others.

As a result, some nodes get overloaded with data while others remain unused. Techniques like data splitting, data shuffling, and the use of sophisticated processing frameworks like Apache Spark can all be used to address data skewness.

[ETL Syllabus](#)

## ETL Interview Questions and Answers for

## Experienced

### **13. In terms of performance, which is preferable: joining data first and then filtering it, or filtering it first and then joining it with other sources?**

- Before joining data with other sources, it is preferable to filter the data.
- Eliminating unneeded data as early in the process as possible is an excellent strategy to boost the performance of the ETL process. It cuts down on the amount of time needed for memory processing, I/O, and/or data transfer.
- As a general rule, fewer rows should be processed and data that never reaches the objective should not be transformed.

### **14. In ETL pipelines, how would you manage incremental data loading?**

Incremental data loading refreshes the new or modified data since the last ETL run, as opposed to reloading the entire dataset.

This can be done by noting the last successful ETL run and utilizing timestamps or metadata to determine the delta changes that have happened subsequently.

Strategies including database triggers, timestamp comparisons, and CDC (Change Data Capture) can be used to detect incremental changes and load them into the target system.

### **15. In ETL workflows, how would you manage exceptions and errors?**

ETL operations need to handle errors and exceptions to ensure data dependability and integrity.

To do this, robust error-handling strategies can be applied

- Adding alerting mechanisms to notify stakeholders about severe failures,
- Logging error details for troubleshooting,

- Setting up checkpoints to resume from the point of failure,
- Retrying methods for transient issues.

Moreover, fault-tolerant processing frameworks like Apache Airflow or Apache NiFi can help with ETL workflow management.

## **16. How would you create an ETL architecture that is both fault-tolerant and scalable?**

When creating a scalable and fault-tolerant ETL architecture, many considerations need to be made, including:

- Implementing distributed processing
- Utilizing fault-tolerant storage systems
- Arranging for horizontal scalability
- Incorporating redundancy and failover techniques.

Using microservices design, cloud-based ETL services, containerization technologies like Docker and Kubernetes, data replication, and backup plans are some of the strategies that can help with building a strong ETL architecture.

## **17. When source schemas alter over time, how would you manage schema evolution in ETL pipelines?**

The practice of managing changes to the structure or schema of the source data over time is known as schema evolution. Among the techniques for controlling schema evolution in ETL pipelines are schema inference, schema mapping, and schema evolution policies.

Furthermore, flexible data models like schema-on-read, data serialization formats like Avro or Parquet, and schema versioning approaches can be used to adapt changes to source schemas without affecting downstream processes.

## **18. In ETL, how are the tables analyzed?**

A cost-based optimizer uses the statistics produced by the 'ANALYZE' statement to determine the most economical strategy for data retrieval. The ANALYZE statement can be used to support the system's object structure validation and space management.

*COMPUTER, ESTIMATE, and DELETE are examples of operations.*

Example Oracle 7 Code:

```
select OWNER,  
  
sum(decode(nvl(NUM_ROWS,9999),  
9999,0,1)) analyzed,  
  
sum(decode(nvl(NUM_ROWS,9999),  
9999,1,0)) not_analyzed,  
  
count(TABLE_NAME) total  
  
from dba_tables  
  
where OWNER not in ('SYS', 'SYSTEM')  
  
group by OWNER
```

To obtain table information, this software runs a SQL query against the database's `dba\_tables` view. The goal is to count the number of analyzed and unanalyzed tables for each owner by analyzing the data in the tables.

## **19. How would you set up the ETL process's logging?**

To monitor all changes and failures during a load, logging is crucial. Using flat files or logging tables is the most popular method for getting ready for logging. In other words, counts, timestamps, and

metadata related to the source and target are added during the operation and subsequently deposited into a table or flat file.

To accomplish that in an ETL process, a developer would utilize variables (such as the Transact-SQL system variable @@ROWCOUNT) or event handlers (SSIS) to track inserted, updated, and deleted entries.

The SSISDB database allows us to monitor each processed package while utilizing SSIS:

```
SELECT TOP 10000 o.Object_Name  
    [Project Name]  
  
    ,REPLACE(e.package_name ,  
    'dtsx', ') [Package Name]  
  
    ,o.start_time                [Start  
    Time]  
  
    ,o.end_time                  [End  
    Time]  
  
    ,e.message_source_name  
    [Message Source Name]  
  
    ,e.event_name                [Event  
    Name]  
  
    ,e.subcomponent_name  
    [Subcomponent Name]  
  
    ,e.message_code  
    [Message Code]  
  
    ,m.message_time
```

*[Event Time]*

*,m.message* *[Error*  
*Message]*

*,m.message\_time*  
*[Error Date]*

*,o.caller\_name* *[Caller*  
*Name]*

*,o.Stopped\_By\_Name*  
*[Stopped By]*

*,ROW\_NUMBER() OVER*  
*(PARTITION BY m.operation\_id ORDER*  
*BY m.message\_source\_type DESC*  
*[Source Type Order]*

*FROM SSISDB.internal.operations o*

*JOIN*  
*SSISDB.internal.operation\_messages m*

*ON o.operation\_id = m.operation\_id*

*JOIN SSISDB.internal.event\_messages*  
*e*

*ON m.operation\_id = e.operation\_id*

*AND m.operation\_message\_id =*  
*e.event\_message\_id*

*WHERE o.Object\_Name LIKE '[object*  
*name]%' – database*



*AND event\_name LIKE '%Error%'*

*ORDER BY o.end\_time DESC*

*,o.Object\_Name*

*,[Source Type Order] ASC*

ETL tools provide native alerting and logging features, such as a dashboard showing the current load status, in addition to flat files and the database itself.

## **20. Could you define OLAP cubes and cubes?**

The cube is an essential component and plays a major part in data processing. Cubes are essentially the Data Warehouse's data processing units; they include dimensions and fact tables. They help clients by providing a multifaceted view of the data in addition to searching and analyzing it.

However, software called Online Analytical Processing (OLAP) makes it possible to analyze data from several databases at once. An OLAP cube can be used to store data in a multidimensional format for reporting purposes.

Cubes simplify the process of creating and reading reports, which improves and streamlines the reporting procedure. End users are in charge of maintaining and managing these cubes, which means they have to manually update the data within.

**ETL Salary**

## **Conclusion**

We hope that our list of ETL interview questions and answers will be helpful to you as you get ready for data science jobs. Enroll in our [ETL Course in Chennai](#) and earn IBM certification that helps you launch your career with in-demand ETL skills.

[ETL Training](#)

Share on your Social Media



## Softlogic Academy

## Softlogic Systems

### KK Nagar [Corporate Office]

No.10, PT Rajan Salai, K.K. Nagar, Chennai  
– 600 078.

**Landmark:** Karnataka Bank Building

**Phone:** [+91 86818 84318](tel:+918681884318)

**Email:** [enquiry@softlogicsys.in](mailto:enquiry@softlogicsys.in)

**Map:** [Google Maps Link](#)

### OMR

No. E1-A10, RTS Food Street  
92, Rajiv Gandhi Salai (OMR),  
Navalur, Chennai – 600 130.

**Landmark:** Adj. to AGS Cinemas

**Phone:** [+91 89256 88858](tel:+918925688858)

**Email:** [info@softlogicsys.in](mailto:info@softlogicsys.in)

**Map:** [Google Maps Link](#)

## Courses

## Navigation

- [About Us](#)
- [Blog Posts](#)
- [Careers](#)
- [Contact](#)
- [Placement Training](#)
- [Corporate Training](#)
- [Hire With Us](#)
- [Job Seekers](#)
- [SLA's Recently Placed Students](#)
- [Reviews](#)
- [Sitemap](#)

## Important Links

- [Disclaimer](#)
- [Privacy Policy](#)
- [Terms and Conditions](#)

## Social Media Links

Python

Software Testing

Full Stack Developer

Java

Power BI

Clinical SAS

Data Science

Embedded


Cloud Computing

Hardware and Networking

VBA Macros

Mobile App Development

DevOps



## Review Sources

Google

Trustpilot

Glassdoor

Mouthshut

Sulekha

Justdial

Ambitionbox

Indeed

Software Suggest

Sitejabber