



# Sybase to Postgres®: A technical migration playbook

What converts automatically, what needs engineering judgment, and how to scope a SAP® Adaptive Server Enterprise migration before committing to a cutover date.

# What this playbook covers

This document is written for the people who will actually run the migration — DBAs, migration architects, and the application owners whose code talks to ASE. It answers three questions in order: what EDB officially supports, what converts without a rewrite, and how long the work takes once you pick a downtime tolerance.

It deliberately does not make the business case. If you need the compliance and audit framing — SAP's maintenance timeline mapped to PCI DSS v4.0.1 and SOX ITGC language — that is a separate companion document, *Migrating from Sybase (SAP ASE): A Compliance Brief*.

# What EDB officially supports

EDB Migration Toolkit lists SAP Adaptive Server Enterprise 15.7 and 16.0 as supported source databases, migrating into EDB Postgres AI (EDB PG AI). This is documented, generally available product support — not a services-only capability.

For ASE sources the toolkit migrates schemas, tables, constraints, indexes, and table data. Views, stored procedures, triggers, and functions are not migrated by the toolkit; they are rewritten as part of the project.

## How this differs from an Oracle or MySQL migration

If you have run one of those before, calibrate carefully — the tooling economics are not the same. Oracle migrations have ora2pg and MySQL has pgloader: mature, free, community-maintained converters with years of accumulated edge-case handling. ASE has no equivalent.

The practical planning consequences are three. First, you cannot assume a free proof-of-concept conversion is available to size the work before budget. Second, community troubleshooting for ASE-specific conversion problems is thin, so budget more internal investigation time per unresolved construct than an Oracle project would need. Third, most available converters come from vendors who do not support the database you land on, which puts the post-cutover support boundary in a different place than you may be used to.

## What the tooling does and does not carry

EDB PG AI is the documented Migration Toolkit target for ASE, so schema and data movement are tooling-assisted. The procedural tier is not.

On a typical ASE estate, stored procedures and triggers are the bulk of the migration effort, and they are rewritten in PL/pgSQL. Plan the project around that rewrite: procedural density shapes the timeline more than any other early decision.



# Construct-by-construct: what converts

Ordered roughly by how much engineering judgment each one demands.

| ASE CONSTRUCT                                        | EDB PG AI EQUIVALENT              | WHAT TO EXPECT                                                                                                                                                                                                 |
|------------------------------------------------------|-----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Tables, indexes, constraints</b>                  | Direct schema objects             | The deterministic part. Migrated by Migration Toolkit with minimal intervention. Check constraints and default clauses need the <code>-ignoreCheckConstFilter</code> option.                                   |
| <b>Views</b>                                         | Postgres views                    | Not migrated by the toolkit for ASE. Recreate them once the base tables land; review any that depend on T-SQL functions.                                                                                       |
| <b>IDENTITY columns</b>                              | Identity columns or sequences     | Confirm in the schema assessment how each IDENTITY column is carried over, then verify seed and increment behavior after cutover.                                                                              |
| <b>T-SQL stored procedures and triggers</b>          | PL/pgSQL, rewritten               | The bulk of the effort on any ASE estate. The toolkit does not migrate procedures, triggers, or functions from ASE, so this is a rewrite, and procedural density is the single biggest driver of timeline.     |
| <b>Temporary tables</b>                              | TEMP tables with Postgres scoping | Semantics differ. Scope and lifetime need review wherever procedures depend on ASE behavior.                                                                                                                   |
| <b>Lock-based concurrency</b>                        | MVCC                              | Not a translation — a behavioral change. Code written around ASE's deterministic locking needs isolation-level review. This is where latency-sensitive systems need real load testing, not functional testing. |
| <b>Application drivers and transaction semantics</b> | PostgreSQL drivers                | Application-layer work, outside the database migration. Scope and staff it separately.                                                                                                                         |

Construct mappings reflect EDB Migration Toolkit's documented object support for ASE sources. Effort per object varies with procedural complexity — confirm against a schema assessment of your own estate.

## The three architectural differences worth planning around

- **Concurrency.** ASE's lock-based model and Postgres's MVCC produce different behavior under contention. Procedures that rely on deterministic lock ordering are the ones to review first.
- **Scaling model.** ASE estates are typically scaled vertically. Postgres adds table partitioning and, with EDB PG AI Distributed High Availability, multi-master replication — which changes the options available for the target architecture, not just the migration.
- **Data types.** ASE has no native vector type. If the reason for moving includes AI workloads, that capability arrives with the target rather than being migrated.



# The five stages

The same five stages, in the same order, on every source database EDB migrates.

| STAGE         | WHAT HAPPENS                                                                                                                                          | EXIT CRITERION                                                                      |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|
| 1 • Assess    | Schema and complexity analysis across the estate; procedural inventory; identification of latency-sensitive and transaction-isolation-dependent code. | A scoped object inventory with effort classification.                               |
| 2 • Transform | Schema and data definitions migrated by Migration Toolkit; procedural code rewritten in PL/pgSQL.                                                     | Converted schema that builds cleanly on the target.                                 |
| 3 • Remediate | Resolution of the edge cases conversion leaves behind — semantics, performance characteristics, application-level expectations.                       | No unresolved conversion exceptions.                                                |
| 4 • Move      | Coordinated bulk data transfer, then change data capture to keep source and target in sync until cutover.                                             | Reconciled row counts and checksums.                                                |
| 5 • Validate  | Integrity, functional correctness, and performance verification under representative load.                                                            | Signed-off test evidence, including audit and regulatory validation where in scope. |

## Scoping note: how these stages are actually executed for ASE

Read this before you size the engagement. EDB's agentic migration capability is an emerging capability, currently scoped to Oracle® and cloud-locked Postgres sources. It is not the delivery mechanism for a Sybase project today.

For ASE, Migration Toolkit carries the schema and bulk-data work in stages 2 and 4. Procedural conversion in stage 2, and all of stages 3 and 5, are predominantly human engineering effort, and stage 3 is where estimates go wrong most often. Plan staffing accordingly: on an ASE project the tooling compresses data movement far more than it compresses code conversion.



# Downtime tolerance drives everything else

Every migration runs two clocks: how long the project takes, and how long the system is offline at cutover. Pick the downtime tolerance first — it determines the strategy, and the strategy determines the timeline.

| STRATEGY           | MECHANISM                  | BEST FOR                            | TARGET DOWNTIME |
|--------------------|----------------------------|-------------------------------------|-----------------|
| Offline            | Backup and restore         | Small, low-complexity databases     | 2–16 hr         |
| Near-zero downtime | CDC / logical replication  | Medium to large databases           | <30–45 min      |
| Phased / shadow    | Parallel run, then cutover | Very large, high-complexity systems | 0 min           |

Downtime figures above are the cross-source envelope for each strategy. Sybase-specific figures are in the next table and are narrower. Source: EDB migration engineering estimates, compiled from completed customer engagements, 2026 (internal; actual ranges vary by estate and are confirmed during assessment).

## Sybase planning ranges, in context

| SOURCE           | OFFLINE                | NEAR-ZERO DOWNTIME         | PHASED / SHADOW        |
|------------------|------------------------|----------------------------|------------------------|
| Sybase / SAP ASE | 3–6 wks · 6–16 hr down | 8–18 wks · <30 min down    | 20–44 wks · 0 min down |
| SQL Server       | 2–4 wks · 4–12 hr down | 6–16 wks · <30 min down    | 16–34 wks · 0 min down |
| DB2              | 3–6 wks · 6–16 hr down | 8–20 wks · <30–45 min down | 26–52 wks · 0 min down |
| MySQL            | 1–2 wks · 2–8 hr down  | 4–12 wks · <15 min down    | 12–26 wks · 0 min down |

Planning ranges, not commitments. Sybase sits at the slower end because the ASE dialect and specialized expertise stretch both assessment and testing. Ranges vary with stored-procedure complexity, custom code, and data volume. Source: EDB migration engineering estimates, compiled from completed customer engagements, 2026 (internal; actual ranges vary by estate and are confirmed during assessment).

# What each stage actually involves

---

## Assessment and workload profiling

Start with the procedural tier, not the schema. Financial and trading systems in particular carry heavy stored-procedure logic, and that inventory — not table count — predicts the timeline. Profile transaction isolation requirements and identify latency-sensitive paths before any conversion begins.

## Schema and code conversion

Migration Toolkit migrates schemas, tables, constraints, indexes, and data. Views, stored procedures, triggers, and functions are rewritten in PL/pgSQL, and that rewrite is where remediation effort concentrates. Confirm how IDENTITY columns land, and expect temp-table semantics to need review.

## Data migration

Parallel bulk transfer, then change data capture for ongoing synchronization until cutover. Migration Toolkit handles the bulk transfer, not ongoing sync, so choose the CDC approach separately. Regulated environments should plan reconciliation checks and audit trails as a first-class workstream rather than a validation afterthought.

## Application remediation

Applications move to PostgreSQL drivers and Postgres transaction semantics. On financial systems this is frequently the longest pole — it is application engineering work, and it is often owned by a different team than the database migration.

## Performance engineering

Memory configuration for high-throughput transactions, index strategy for the query patterns that matter, and partitioning for large tables. Do this against representative load, because the MVCC transition means functional parity does not imply performance parity.

## Cutover

Executed in a controlled window with a tested rollback path. Dual-run or shadow environments are standard practice where the system cannot tolerate a failed cutover.



# Where to start

## 1

### Inventory the procedural tier

Count and classify stored procedures, triggers, and the application code paths that depend on ASE-specific transaction behavior. This is the number that drives every estimate that follows.

## 2

### Decide your downtime tolerance

Agree it with the business before scoping. It selects the strategy, and the strategy sets the timeline band from the table in section 05.

## 3

### Run a schema assessment

Scope the conversion against your real estate rather than a planning range. Assessment output is also the input a compliance remediation plan needs if ASE 16.0 is in PCI or SOX scope.

## 4

### Review findings with a migration architect

Bring the inventory and the assessment. The output is a scoped timeline suitable for a budget conversation — not a commitment made before the work is understood.

### Request an assessment

Go to [enterprisedb.com/migration/sybase](https://enterprisedb.com/migration/sybase), or contact your EDB account team. If you received this document from an EDB representative, reply to them directly — they can schedule the assessment review in step 4.

[enterprisedb.com/migration/sybase](https://enterprisedb.com/migration/sybase)

## Sources

- EDB Migration Toolkit documentation — supported platforms and databases ([enterprisedb.com/docs/migration\\_toolkit/latest/02\\_supported\\_operating\\_systems\\_and\\_database\\_versions](https://enterprisedb.com/docs/migration_toolkit/latest/02_supported_operating_systems_and_database_versions)).
- EDB Migration Toolkit documentation — overview and object migration support by source database ([enterprisedb.com/docs/migration\\_toolkit/latest/04\\_functionality\\_overview](https://enterprisedb.com/docs/migration_toolkit/latest/04_functionality_overview)).
- EDB migration engineering estimates, compiled from completed customer engagements, 2026 (internal; actual ranges vary by estate and are confirmed during assessment).
- SAP Note 1922006, EoMM Dates for Adaptive Server Enterprise 15.0.3, 15.5, 15.7, 16.0 and 16.1 — SAP ASE 16.0 end of mainstream maintenance December 31, 2027; SAP ASE 16.1 end of mainstream maintenance December 31, 2030.



# Confirm scope against an assessment of your own estate.

---

Planning ranges and construct mappings in this document are engineering guidance, not contractual commitments.

This playbook is engineering planning guidance and does not constitute a contractual commitment, a performance guarantee, or legal, audit, or compliance advice. Timeline ranges are planning estimates that vary by environment. SAP, Sybase, and SAP Adaptive Server Enterprise are trademarks of SAP SE. Oracle is a registered trademark of Oracle Corporation. Postgres and PostgreSQL are registered trademarks of the PostgreSQL Community Association of Canada and used with their permission. EDB Postgres® AI is a trademark of EnterpriseDB Corporation. © EnterpriseDB Corporation 2026. All rights reserved.

EDB Postgres® AI is the first open, enterprise-grade sovereign data and AI platform, with a secure, compliant, and fully scalable environment, on premises and across clouds. Supported by a global partner network, EDB Postgres AI unifies transactional, analytical, and AI workloads, enabling organizations to operationalize their data and LLMs where, when, and how they need them. For more information, visit [enterprisedb.com](https://enterprisedb.com).

