

Indexing Strategy

Indexing Strategy

Version: 1.0 **Date:** 2026-02-21 **Status:** Approved **Owner:** Database Architect

Overview

Drop's indexing strategy is designed around actual user flow query patterns. Every index exists because a specific query needs it. No speculative indexes.

Current index count: 16 indexes across 19 tables (defined in `db.ts` schema).

Query Patterns by User Flow

Login Flow (BankID OIDC)

```
sequenceDiagram
    participant U as User
    participant API as API
    participant DB as Database

    U->>API: BankID callback (code, state)
    API->>DB: Q1: Find user by national_id_hash
    DB-->>API: User or NULL
    alt New user
        API->>DB: Q2: INSERT users
        API->>DB: Q3: INSERT settings (defaults)
    end
    API->>DB: Q4: INSERT sessions
    API-->>U: JWT cookie + redirect
```

Query ID	SQL Pattern	Index Required	Current Coverage
Q1	SELECT * FROM users WHERE national_id_hash = ?	idx_users_national_id (partial: WHERE NOT NULL)	Covered
Q2	INSERT INTO users (...)	None (PK insert)	N/A
Q3	INSERT INTO settings (...)	None (PK insert)	N/A
Q4	INSERT INTO sessions (...)	None (PK insert)	N/A

Authentication Middleware (every authenticated request)

Query ID	SQL Pattern	Index Required	Current Coverage
Q5	SELECT * FROM sessions WHERE token_hash = ? AND revoked = 0 AND expires_at > ?	idx_sessions_token	Covered
Q6	SELECT * FROM users WHERE id = ?	PRIMARY KEY	Covered

Dashboard View

```
sequenceDiagram
    participant U as User
    participant API as API
    participant DB as Database

    U->>API: GET /auth/me
    API->>DB: Q5: Verify session (token_hash)
    API->>DB: Q6: Get user by PK
    API->>DB: Q7: Get bank accounts for user
    DB-->>API: Bank accounts with cached balances
    API-->>U: User profile + total balance
```

Query ID	SQL Pattern	Index Required	Current Coverage
Q7	SELECT * FROM bank_accounts WHERE user_id = ?	idx_bank_accounts_user	Covered

Transaction History

Query ID	SQL Pattern	Index Required	Current Coverage
Q8	SELECT * FROM transactions WHERE user_id = ? [AND type = ?] [AND status = ?] ORDER BY created_at DESC LIMIT ? OFFSET ?	idx_transactions_user	Covered (user_id)
Q9	SELECT COUNT(*) FROM transactions WHERE user_id = ?	idx_transactions_user	Covered

Note: The `type` and `status` filters are applied after the `user_id` index lookup. At current scale (< 10K transactions per user), this is efficient. A composite index `(user_id, created_at DESC)` would optimize the ORDER BY for users with many transactions.

Create Remittance

Query ID	SQL Pattern	Index Required	Current Coverage
Q10	SELECT * FROM recipients WHERE id = ? AND user_id = ?	idx_recipients_user + PK	Covered
Q11	SELECT * FROM exchange_rates WHERE to_currency = ?	Sequential scan (6 rows)	Acceptable (tiny table)
Q12	SELECT * FROM bank_accounts WHERE user_id = ? AND is_primary = 1	idx_bank_accounts_user	Covered
Q13	UPDATE bank_accounts SET balance = balance - ? WHERE id = ? AND balance >= ?	PK	Covered
Q14	INSERT INTO transactions (...)	None (PK insert)	N/A

Create QR Payment

Query ID	SQL Pattern	Index Required	Current Coverage
Q15	SELECT * FROM merchants WHERE id = ?	PK	Covered
Q12	(Same as remittance -- primary bank account)	idx_bank_accounts_user	Covered
Q13	(Same as remittance -- balance debit)	PK	Covered
Q14	(Same as remittance -- insert transaction)	N/A	N/A

Notifications List

Query ID	SQL Pattern	Index Required	Current Coverage
Q16	<code>SELECT * FROM notifications WHERE user_id = ? ORDER BY created_at DESC</code>	<code>idx_notifications_user</code>	Covered
Q17	<code>UPDATE notifications SET read = 1 WHERE id IN (?, ?, ...) AND user_id = ?</code>	PK + <code>idx_notifications_user</code>	Covered

Settings View/Update

Query ID	SQL Pattern	Index Required	Current Coverage
Q18	<code>SELECT * FROM settings WHERE user_id = ?</code>	PK (user_id IS the PK)	Covered
Q19	<code>UPDATE settings SET ... WHERE user_id = ?</code>	PK	Covered

Merchant Dashboard

Query ID	SQL Pattern	Index Required	Current Coverage
Q20	<code>SELECT * FROM merchants WHERE user_id = ?</code>	Sequential scan (1 merchant per user)	Acceptable
Q21	<code>SELECT * FROM transactions WHERE merchant_id = ? AND created_at >= ?</code>	<code>idx_transactions_merchant</code>	Partially covered (no composite with created_at)

Recipient Management

Query ID	SQL Pattern	Index Required	Current Coverage
Q22	<code>SELECT * FROM recipients WHERE user_id = ? LIMIT ? OFFSET ?</code>	<code>idx_recipients_user</code>	Covered
Q23	<code>SELECT * FROM recipients WHERE id = ? AND user_id = ?</code>	PK + <code>idx_recipients_user</code>	Covered
Q24	<code>DELETE FROM recipients WHERE id = ? AND user_id = ?</code>	PK	Covered

Compliance Queries (Admin/Internal)

Query ID	SQL Pattern	Index Required	Current Coverage
----------	-------------	----------------	------------------

Q25	<code>SELECT * FROM audit_log WHERE user_id = ? ORDER BY timestamp DESC</code>	<code>idx_audit_log_user</code>	Covered
Q26	<code>SELECT * FROM audit_log WHERE action = ? AND timestamp BETWEEN ? AND ?</code>	<code>idx_audit_log_action + idx_audit_log_timestamp</code>	Partially (separate indexes, no composite)
Q27	<code>SELECT * FROM aml_alerts WHERE user_id = ? AND status IN ('open', 'investigating')</code>	<code>idx_aml_alerts_user</code>	Covered
Q28	<code>SELECT * FROM aml_alerts WHERE status = 'open' ORDER BY created_at</code>	<code>idx_aml_alerts_status (proposed)</code>	Not covered (needs new index)
Q29	<code>SELECT * FROM complaints WHERE user_id = ? ORDER BY created_at DESC</code>	<code>idx_complaints_user</code>	Covered
Q30	<code>SELECT * FROM complaints WHERE status IN ('received', 'investigating')</code>	<code>idx_complaints_status (proposed)</code>	Not covered (needs new index)

Index Inventory

Current Indexes (defined in `db.ts`)

Index Name	Table	Column(s)	Type	Rationale
<code>idx_users_national_id</code>	<code>users</code>	<code>national_id_hash</code>	B-tree, partial (WHERE NOT NULL)	BankID login deduplication -- find user by hashed national ID
<code>idx_recipients_user</code>	<code>recipients</code>	<code>user_id</code>	B-tree	List recipients per user, verify ownership
<code>idx_transactions_user</code>	<code>transactions</code>	<code>user_id</code>	B-tree	Transaction history per user (most frequent query)
<code>idx_transactions_merchant</code>	<code>transactions</code>	<code>merchant_id</code>	B-tree	Merchant dashboard -- transactions for merchant (documented in DATABASE- SCHEMA.md)
<code>idx_tx_idempotency</code>	<code>transactions</code>	<code>idempotency_key</code>	B-tree, unique, partial (WHERE NOT NULL)	Prevent duplicate transaction submission

Index Name	Table	Column(s)	Type	Rationale
<code>idx_bank_accounts_user</code>	<code>bank_accounts</code>	<code>user_id</code>	B-tree	Dashboard balance lookup, transaction source
<code>idx_sessions_user</code>	<code>sessions</code>	<code>user_id</code>	B-tree	Revoke all sessions on logout
<code>idx_sessions_token</code>	<code>sessions</code>	<code>token_hash</code>	B-tree	Auth middleware -- validate session on every request
<code>idx_notifications_user</code>	<code>notifications</code>	<code>user_id</code>	B-tree	Notifications list per user (documented in DATABASE-SCHEMA.md)
<code>idx_audit_log_user</code>	<code>audit_log</code>	<code>user_id</code>	B-tree	User investigation, DSAR compliance
<code>idx_audit_log_action</code>	<code>audit_log</code>	<code>action</code>	B-tree	Event type filtering for monitoring
<code>idx_audit_log_timestamp</code>	<code>audit_log</code>	<code>timestamp</code>	B-tree	Time-range queries for compliance reporting (documented in DATABASE-SCHEMA.md)
<code>idx_aml_alerts_user</code>	<code>aml_alerts</code>	<code>user_id</code>	B-tree	Per-user AML alert lookup
<code>idx_aml_alerts_status</code>	<code>aml_alerts</code>	<code>status</code>	B-tree	Open alerts dashboard (documented in DATABASE-SCHEMA.md)
<code>idx_complaints_user</code>	<code>complaints</code>	<code>user_id</code>	B-tree	Per-user complaint history
<code>idx_screening_user</code>	<code>screening_results</code>	<code>user_id</code>	B-tree	Per-user screening history

Indexes from DATABASE-SCHEMA.md (not in db.ts code)

The DATABASE-SCHEMA.md documentation lists additional indexes that may not be in the current `db.ts` `SQLITE_SCHEMA` string:

Index Name	Table	Column(s)	Status
------------	-------	-----------	--------

<code>idx_merchants_org</code>	<code>merchants</code>	<code>org_number</code>	Documented but covered by UNIQUE constraint
<code>idx_cards_user</code>	<code>cards</code>	<code>user_id</code>	Documented, may not be in db.ts
<code>idx_spending_limits_user</code>	<code>spending_limits</code>	<code>user_id</code>	Documented, may not be in db.ts
<code>idx_spending_limits_card</code>	<code>spending_limits</code>	<code>card_id</code>	Documented, may not be in db.ts
<code>idx_consents_user</code>	<code>consents</code>	<code>user_id</code>	Documented, may not be in db.ts
<code>idx_data_requests_user</code>	<code>data_access_requests</code>	<code>user_id</code>	Documented, may not be in db.ts
<code>idx_complaints_status</code>	<code>complaints</code>	<code>status</code>	Documented, may not be in db.ts

Recommendation: Reconcile DATABASE-SCHEMA.md with actual `db.ts` code. Add missing indexes to the schema if the queries justify them.

Recommended Additional Indexes

Based on query pattern analysis, the following indexes should be added:

Proposed Index	Table	Column(s)	Justification
<code>idx_transactions_user_created</code>	<code>transactions</code>	<code>(user_id, created_at DESC)</code>	Optimizes paginated transaction history (Q8) -- avoids sort after index lookup
<code>idx_complaints_status</code>	<code>complaints</code>	<code>status</code>	Admin dashboard query for open complaints (Q30)
<code>idx_consents_user</code>	<code>consents</code>	<code>user_id</code>	DSAR export needs all consents for user
<code>idx_data_requests_user</code>	<code>data_access_requests</code>	<code>user_id</code>	DSAR tracking per user
<code>idx_audit_log_resource</code>	<code>audit_log</code>	<code>(resource_type, resource_id)</code>	Resource-specific audit trail lookup

Partial Index Opportunities (PostgreSQL)

These are PostgreSQL-specific optimizations to add after migration:

Proposed Index	Table	Column(s)	Condition	Justification
----------------	-------	-----------	-----------	---------------

<code>idx_sessions_active</code>	<code>sessions</code>	<code>user_id</code>	<code>WHERE revoked = 0</code>	Auth middleware only queries active sessions
<code>idx_aml_alerts_open</code>	<code>aml_alerts</code>	<code>created_at</code>	<code>WHERE status IN ('open', 'investigating')</code>	Dashboard shows only open alerts
<code>idx_notifications_unread</code>	<code>notifications</code>	<code>user_id</code>	<code>WHERE read = 0</code>	Badge count for unread notifications
<code>idx_users_active</code>	<code>users</code>	<code>email</code>	<code>WHERE deleted_at IS NULL</code>	Login only checks non-deleted users

Connection Pooling Configuration

“ **Note (ADR-014, 2026-03-03):** Drop uses PostgreSQL 16 in ALL environments. SQLite and the dual-driver layer have been removed. The section below reflects the current PostgreSQL-only configuration.

PostgreSQL 16 (All Environments)

PostgreSQL uses Drizzle ORM with connection pooling:

Setting	Value	Source	Notes
Pool library	<code>pg.Pool</code>	<code>db.ts:16-21</code>	Node-postgres built-in pool
Connection string	<code>DATABASE_URL</code> env var	<code>db.ts:18</code>	Standard PostgreSQL URL format
Max connections	Default (10)	<code>pg.Pool</code> default	Adjust based on App Runner instance count
Idle timeout	10,000ms	<code>pg.Pool</code> default	Close idle connections after 10s
Connection timeout	0 (no timeout)	<code>pg.Pool</code> default	Wait indefinitely for connection

Recommended Production Pool Configuration

```
// Recommended pg.Pool configuration for production
const pool = new pg.Pool({
  connectionString: process.env.DATABASE_URL,
```

```

max: 20, // Max connections per instance
idleTimeoutMillis: 30000, // Close idle after 30s
connectionTimeoutMillis: 5000, // Fail if no connection in 5s
ssl: { rejectUnauthorized: true } // Require SSL for RDS
});

```

Parameter	Recommended Value	Rationale
max	20	Balance between connection availability and RDS connection limits. With 2-3 App Runner instances, total connections = 40-60 (well under RDS default 100).
idleTimeoutMillis	30,000	Close idle connections to free RDS slots, but keep them long enough to avoid reconnection overhead for bursty traffic.
connectionTimeoutMillis	5,000	Fail fast on connection issues rather than hanging. API should return 503 to client.
ssl	{ rejectUnauthorized: true }	Encrypt connections to RDS. Required for compliance.

PgBouncer Consideration

At current projected scale (3,000 users, ~100 concurrent connections), direct `pg.Pool` is sufficient. PgBouncer should be evaluated when:

- Connection count exceeds RDS limits
- Multiple services need to share the same database
- Transaction-mode pooling would reduce connection overhead

EXPLAIN ANALYZE Examples

Transaction History Query (most common)

```

-- PostgreSQL 16 (all environments - ADR-014)
EXPLAIN ANALYZE
SELECT * FROM transactions
WHERE user_id = 'usr_demo1'

```

```
ORDER BY created_at DESC
LIMIT 20 OFFSET 0;
-- Expected: Index Scan using idx_transactions_user on transactions
--          Sort Key: created_at DESC
--          Rows Removed by Index: 0 (all rows match user_id)
```

Session Validation (every request)

```
-- PostgreSQL 16
EXPLAIN ANALYZE
SELECT * FROM sessions
WHERE token_hash = 'abc123...'
AND revoked = FALSE
AND expires_at > NOW();
-- Expected: Index Scan using idx_sessions_token on sessions (token_hash=?)

-- PostgreSQL
EXPLAIN ANALYZE
SELECT * FROM sessions
WHERE token_hash = 'abc123...'
AND revoked = 0
AND expires_at > '2026-02-21T00:00:00';
-- Expected: Index Scan using idx_sessions_token on sessions (cost=0.28..8.30)
--          Filter: (revoked = 0 AND expires_at > ...)
```

Audit Log by User (investigation)

```
-- PostgreSQL
EXPLAIN ANALYZE
SELECT * FROM audit_log
WHERE user_id = 'usr_demo1'
ORDER BY timestamp DESC
LIMIT 100;
-- Expected: Index Scan Backward using idx_audit_log_user
```

Performance Monitoring

Key Metrics to Track

Metric	Target	Alert Threshold	Query
Session validation latency	< 5ms	> 20ms	<code>SELECT * FROM sessions WHERE token_hash = ?</code>
Transaction list latency	< 50ms	> 200ms	<code>SELECT * FROM transactions WHERE user_id = ? ORDER BY created_at DESC LIMIT 20</code>
Audit log write latency	< 10ms	> 50ms	<code>INSERT INTO audit_log (...)</code>
Index bloat	< 20%	> 50%	<code>pg_stat_user_indexes</code>
Sequential scans on large tables	0	Any	<code>pg_stat_user_tables.seq_sca n</code> for transactions, audit_log

Periodic Index Maintenance (PostgreSQL)

```
-- Check index usage
SELECT schemaname, tablename, indexname, idx_scan, idx_tup_read
FROM pg_stat_user_indexes
ORDER BY idx_scan ASC;

-- Check for bloated indexes
SELECT pg_size_pretty(pg_relation_size(indexrelid)) as size, indexrelid::regclass
FROM pg_stat_user_indexes
ORDER BY pg_relation_size(indexrelid) DESC;

-- Reindex if bloated
REINDEX INDEX CONCURRENTLY idx_transactions_user;
```

Cross-References

- **Database schema:** [DATABASE-SCHEMA.md](#)
- **Database design:** [database-design.md](#)
- **Audit architecture:** [audit-architecture.md](#)
- **Data architecture:** [data-architecture.md](#)
- **Migration strategy:** [migration-strategy.md](#) (PostgreSQL-specific optimizations)
- **Drizzle ORM schema:** `src/shared/db/schema.ts`

Revision #8

Created 2026-02-21 05:59:07 UTC by John

Updated 2026-05-23 10:57:21 UTC by John