

Salesforce Data Management

User License Types

User License types control data access & storage.

Internal User Licenses (non-exhaustive):

- Sales Cloud - Sales Objects, Custom Objects (CRUD)
- Service Cloud - Sales Cloud plus extended Service Objects (CRUD).
- Salesforce Platform - Account+Contact (RU), Custom Objects (RU)
- Lightning Platform App Subscription Light - 10 Custom Objects (RU)
- Lightning Platform App Subscription Enterprise - Account+Contact (RU), 10 Custom Objects (RU)
- Employee Community - 10 Custom Objects (RU)

External User Licenses (non-exhaustive):

- Partner Portal - Below plus Opportunity C
- Customer Community Plus - Below plus Account C
- Service Cloud Portal - Account (RU), Contacts (CRU), Assets (CRU), Cases (CRU)
- External Identity - Custom Object (CRUD)

C = Create
R = Read
U = Write
D = Delete

Data Classification

Data Classification Metadata Fields:

- Data Compliance Category - Picklist (PII, GDPR, HIPAA etc.)
- Data Sensitivity - Picklist (Public, Internal, Confidential)
- ... metadata values can be uploaded per Object-Field (E.g. Contact, BirthDate, Confidential)
- Data Owner - User or Public Group responsible for setting sensitivity etc.
- Field Usage - Active, DeprecateCandidate, Hidden

Data classification enables recording of the data owner, field usage, data sensitivity level, and compliance categorisation of fields in any standard or custom object. Data classification can be used to control access, reporting and to manage data compliance.

Data Governance

Data governance is the process of managing the availability, usability, integrity and security of enterprise data. A governance framework must define standards, controls and processes that control data usage and data protection compliance.

Data Protection GDPR:

- Get Buy-in and build your team - leadership buy-in, core cross-functional team to work on compliance effort. Appoint Data Protection officer.
- Assess your organisation - Data inventory for personal data and register of processing activities. Establish controls and processes - create a roadmap of operational and technical changes.
- Privacy notices - Website, cookies, tags etc.
- Usage limitations - Administrative and technological controls to limit use of data.
- Security - Administrative, physical and technological security measures.
- Data Subject rights - process to manage consent preferences and complaints.
- Vendor management - contracts with 3rd parties.
- Incident response - detect and respond to security breaches.
- Training - employee and vendor training for awareness of policy.
- Assessments - Data protection impact assessments for high risk data processing activity.
- Document Compliance - the protection of personal data is a fundamental right.

Enterprise Data Governance Framework:

- Roles & Responsibilities -
- Data Steward - Defines teams, roles and responsibilities for quality improvement and day-to-day maintenance. Quality Control.
- Data Custodian - Responsible for the technical environment.
- Governance Plan - Data Definitions, Data Quality Standards, Permissions, Quality Control.
- Governance Framework Models -
- Centralised - execution of rules, standards and policies by a central governance body.
- Decentralised - individual responsibility for data quality maintenance.
- Hybrid - combines both.

Large Data Volumes

Large Data Volume considerations apply to objects with millions of records.

Query Optimiser: Determines best index for filters, or Object (Table Scan); order for remaining Objects, injects custom FK value objects for efficient join paths, optimises execution plan & updates statistics.

Database Statistics: updated nightly. The Query Optimiser maintains a table containing statistics about the distribution of data in each index.

Indexes: Standard indexes on relationship fields, RecordTypeId, SystemModstamp, Id and Name fields. Custom indexes can be added; excl. multi-select picklist, long text, non-deterministic formula, encrypted fields. External fields are indexed.

For Custom Fields Salesforce creates an Index Table with strongly typed fields and a Database Index. Records with nulls (empty-value) are omitted by default.

Index Usage: internal statistics first checked for selectivity for the query filters.

- Standard Indexes - filter matches less than 30% of the first million records and less than 15% of additional records.
- Custom Indexes - filter matches less than 10% of the total records, up to a maximum of 333,333 records.

WHERE clause contains:

- AND - used unless one index exceeds 20% total or 666K.
- OR - used unless one index exceeds 10% total or 333K. All fields must be indexed.
- LIKE - internal statistics are not used, 100K sample.

Two-Column Indexes: List views, 1 to display, 1 to sort.

Divisions: data partitioning strategy (EMEA, APAC etc.)

Skinny Tables: automatically improve the performance of read-only operations by removing the joins between underlying database tables that store standard and custom fields separately. Useful above Million record threshold. Copied to Full Copy sandboxes only. Also termed "Reporting Accelerators".

Skinny tables can be created on custom objects, and on Account, Contact, Opportunity, Lead, and Case objects. They can enhance performance for reports, list views, and SOQL. Literal Data Types (+encrypted) are supported. Skinny tables can't contain fields from other objects.

Object: Standard (Field1, Field2) vs Custom (Field3, Field4) vs Skinny (Field1, Field2, Field3, Field4, <=100)

Techniques

Mashups: quick, loosely coupled integration of the two applications. Data is never stale. Short interactions, small amounts of data.

Defer Sharing Calculation: suspend and resume sharing calculations, manage two processes: group membership calculation and sharing rule calculation until new users, rules, and other content have been loaded.

SOQL and SOSL:

- SOQL - SOQL==objects+fields are known; single object or related; number, date, checkbox fields; count the records for criteria, sort the results.
- SOSL - SOQL==objects+fields unknown; multiple objects unrelated; efficient division queries. Faster for CONTAINS queries. SOSL can tokenize multiple terms within a field (for example, multiple words separated by spaces) and build a search index.

Deleting Data: Soft delete moves data to the Recycle Bin for 15 days, the data remains resident and impacts on performance. Use Hard Delete option in the Bulk API.

Search: large data volume operations require the search index to be rebuilt.

Best Practice

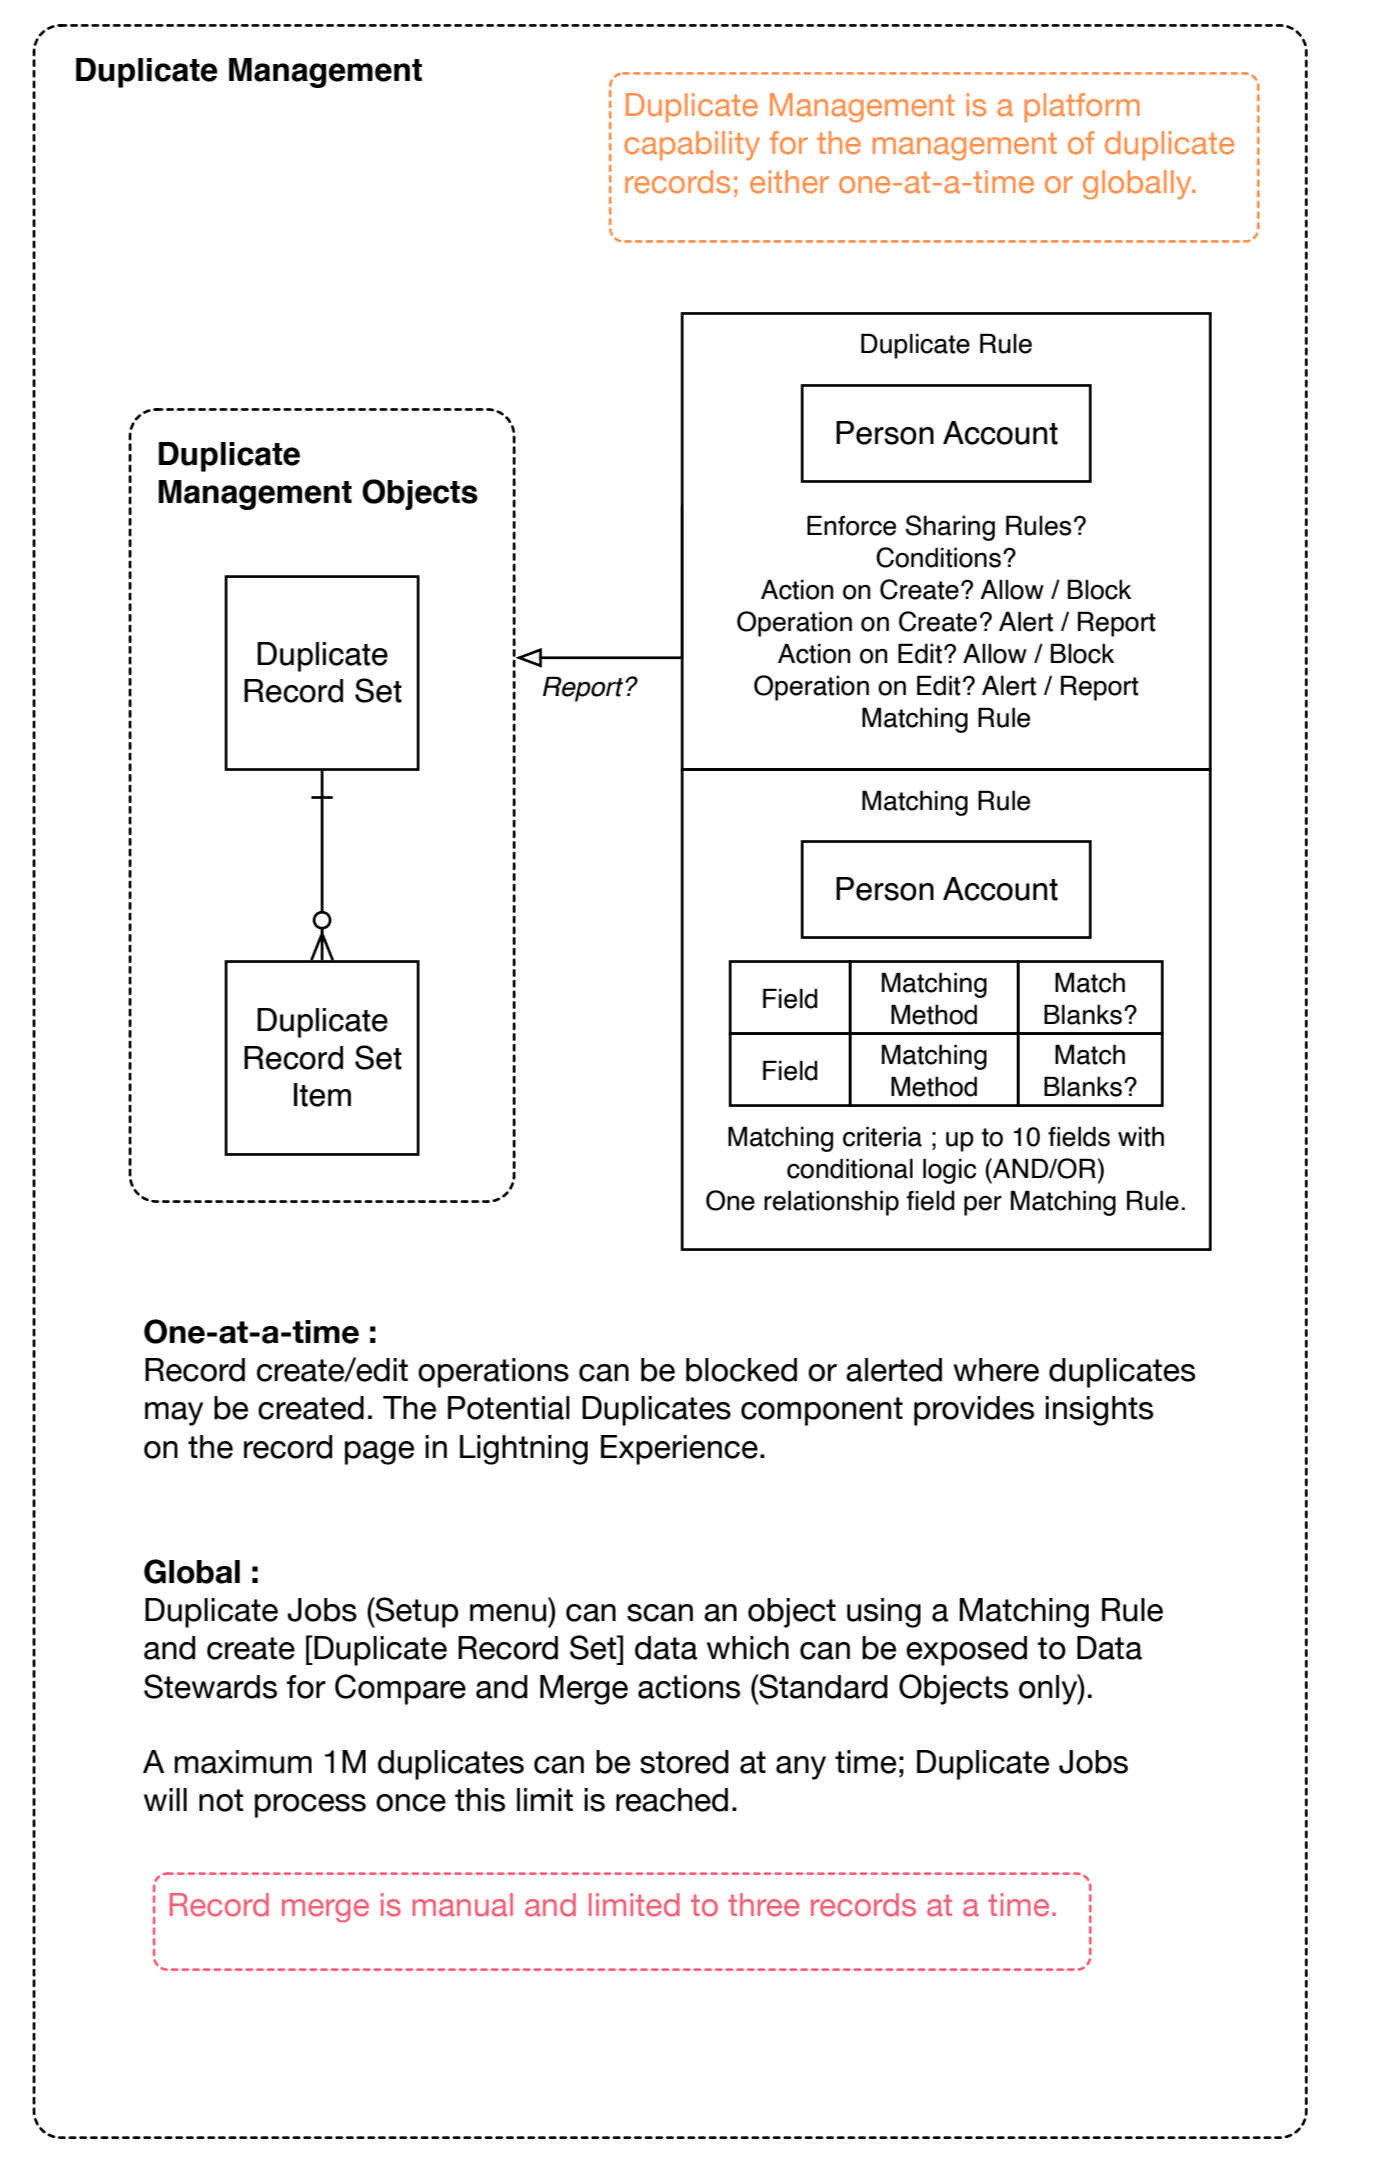
Reporting: Limit relationships, fields and records. Partition data to match its likely use. Reduce records per object. Reduce joins by restricting number of relationships per report. De-normalise data when practical. Limit the number of fields in a report, SOQL or list view to the essential. User filters which emphasise the use of Standard then Custom indexes.

Searching: Keep searches single-object, specific and avoid wildcards.

SOQL and SOSL: Avoid querying of formula fields especially deterministic formulae. Tune the SOQL query, reducing query scope, and using selective filters. Single-column indexes for picklists and foreign key fields exclude rows in which the index column is equal to null. Use a Two-Column index.

Deleting Data: Use Bulk API hard delete option, delete child records first.

General: Avoid having any user own more than 10,000 records. Avoid any parent record having more than 10,000 child records.



Data Import/Export

The Salesforce platform provides a rich set of APIs and tools for the import and export of data in small synchronous batches or via large volume asynchronous data loading operations.

The Bulk API is intended for large volume Data Import, Export and Delete operations.

Loading Data from the API:

- More than 100K use Bulk API.
- Clean data first as errors in batches forces single row processing for the batch.
- Use read/write sharing for initial load.
- If possible for initial loads, populate roles before sharing rules.
- Load users into roles.
- Load record data with owners, triggering role hierarchy calculations.
- Configure public groups & queues, let computations propagate.
- Add sharing rules one at a time, let computations finish for each.
- If possible, add users & data then create & assign groups & queues.
- Minimise parent record-locking conflicts by grouping records by parent.
- Defer Computations (disable ApexTriggers and process automations).
- Defer sharing calculations.
- Tune the batch size (HTTP keepalives, GZIP compression)

Export Data via the Bulk API Query:

- Use the PK Chunking request header (sfce-Enable-PKChunking) to enable automatic primary key (PK) chunking for a bulk query job. PK chunking splits bulk queries on very large objects into chunks based on the record IDs, or primary keys, of the queried records.

Delete Data via the Bulk API:

- For large volume deletes use the Hard Delete option to skip the Recycle Bin (soft-delete) processing.

Cross Org Data Synchronisation:

- Salesforce to Salesforce - Native platform capability. Object Publication/Subscription model.
- Organisation Sync (retired) - A second Salesforce Org with a synchronised sub-set of data.
- Salesforce Connect Cross Org Adapter - External Object access (Virtualised) to data held in a second Salesforce org (Publication org). Consumes REST API calls.