Data Quality Sense

Salesforce-native Data Quality Management

Documentation • April 2026

Table of Contents

Getting Started

Introduction

⚠Try the Beta!

Data Quality Sense is not yet listed on AppExchange. Install the current beta version and try it out on your org!

Install Beta Package Install Data Quality Sense directly on your Salesforce org. →

After installation, the app requires activation. Contact us at hello@tucario.com or see the Installation guide for details.

We'd love to hear your thoughts! Share your feedback.

What is Data Quality Sense?

🔗Section titled "What is Data Quality Sense?"

Data Quality Sense (DQS) is a **Salesforce-native managed package** for monitoring, measuring, and improving the quality of your CRM data. It runs entirely within your Salesforce org — no external integrations, no data leaving your environment.

The standard Salesforce platform doesn't provide built-in tools for measuring data quality at scale. DQS fills this gap with automated scanning, configurable quality dimensions, and rich analytics.

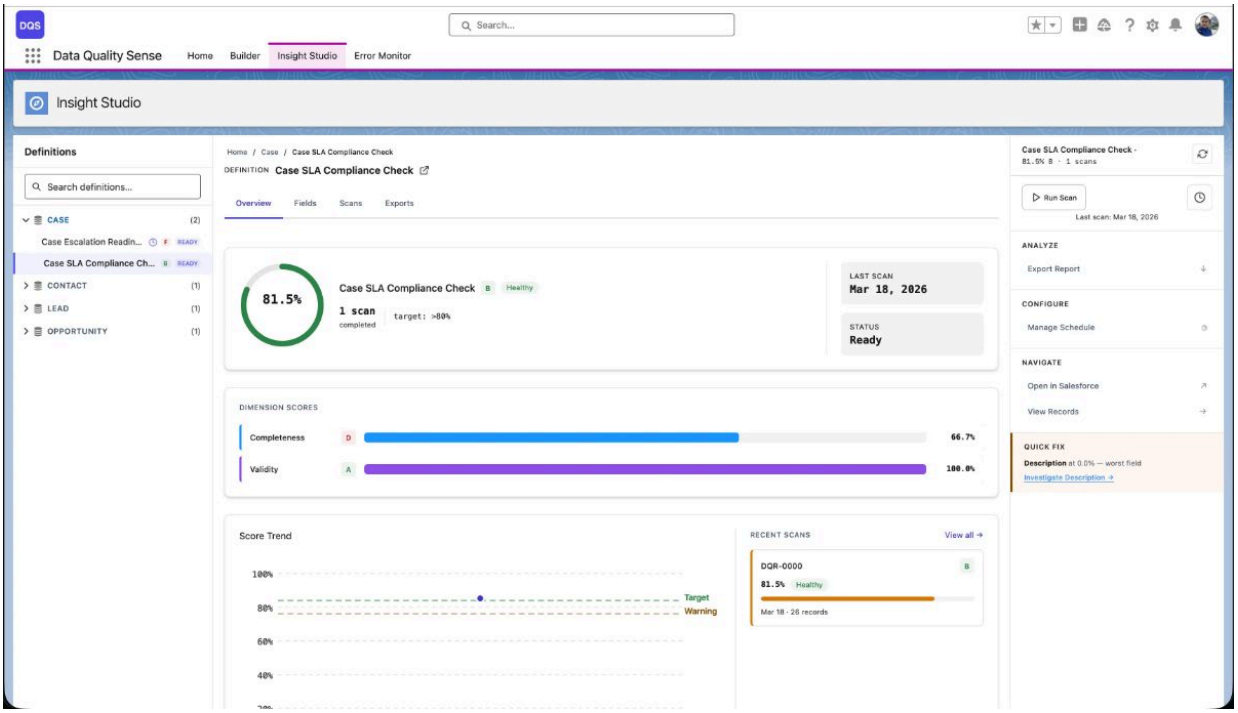

Core Components

Section titled “Core Components”

DQS Builder

Section titled “DQS Builder”

A multi-step wizard for creating scan configurations. Select objects, pick fields, choose quality dimensions, and set thresholds — all through a guided UI.

Insight Studio

Section titled “Insight Studio”

Dashboards and analytics for scan results. Scores, trends, field health matrices, scan comparison, CSV exports, and AI-powered recommendations.

Processing Engine

Section titled “Processing Engine”

Automated batch scanning with flexible scheduling, data retention policies, error management, and platform event notifications.

Quality Dimensions

Section titled “Quality Dimensions”

DQS evaluates data across **6 quality dimensions**:

Dimension	What It Measures
Completeness	Are fields populated?
Validity	Do values match expected formats?
Uniqueness	Are there duplicate values?
Timeliness	Is data up to date?
Consistency	Are related fields logically consistent?
PII Detection	Is personal data properly handled?

Each dimension can be configured globally or overridden per field, with configurable thresholds and scoring.

Key Features

Section titled “Key Features”

- **100% Salesforce-native** — no external dependencies or callouts
- **Any SObject** — scan standard and custom objects
- **Automated scheduling** — CRON-based recurring scans
- **Data retention** — configurable auto-purge of old results
- **CSV export** — download violation details for any dimension
- **Error management** — dedicated console for monitoring scan issues
- **Managed package** — installs cleanly with `dataqualitysens` namespace

Next Steps

Section titled “Next Steps”

- Install the package
- Configure permissions
- Create your first scan

Installation

Try the Beta

Data Quality Sense is not yet listed on AppExchange. You can install the current beta version directly:

Install Beta Package

Prerequisites

Section titled “Prerequisites”

- Salesforce org (Enterprise, Unlimited, or Developer Edition)
- System Administrator profile or equivalent permissions
- Lightning Experience enabled

Install the Package

Section titled “Install the Package”

1. **Get the installation link**

   Contact us at dataqualitysense.com to receive the managed package installation URL for your org.

2. **Choose installation scope**

   Select who should have access:

   - **Install for Admins Only** — recommended for initial setup
   - **Install for All Users** — if all users need immediate access
   - **Install for Specific Profiles** — for granular control

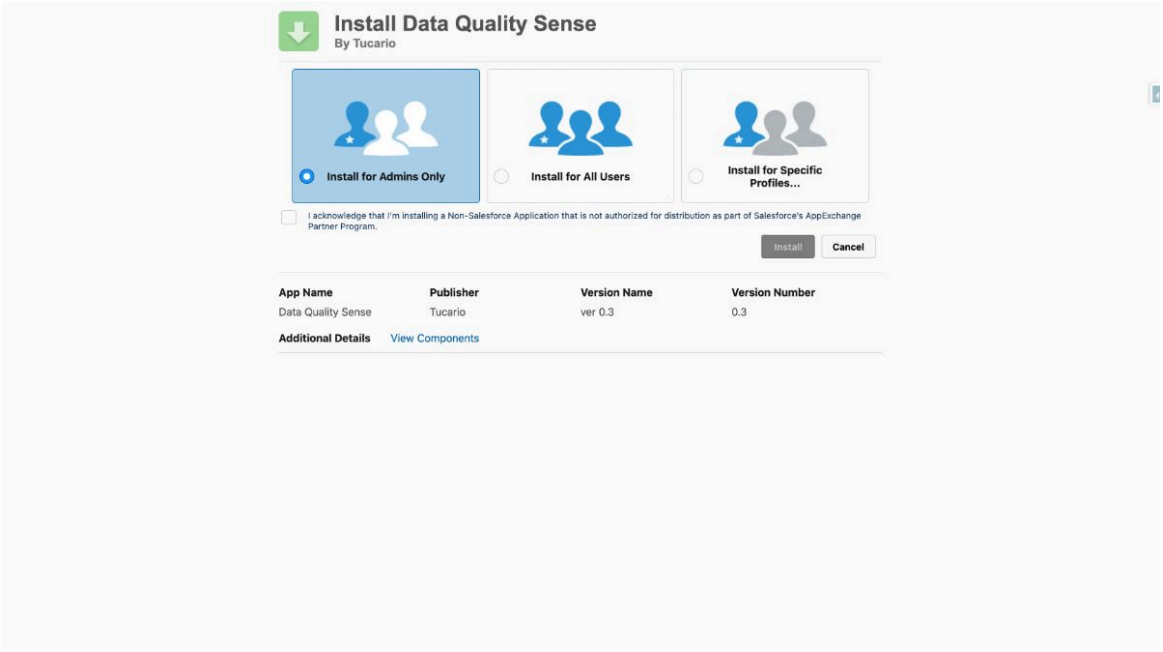

3. **Approve third-party access**

   The package does not require any external callouts. All processing happens within your Salesforce org.

4. **Wait for installation to complete**

The installer will show progress. For larger orgs, installation may take a few minutes — you will receive an email once it completes.

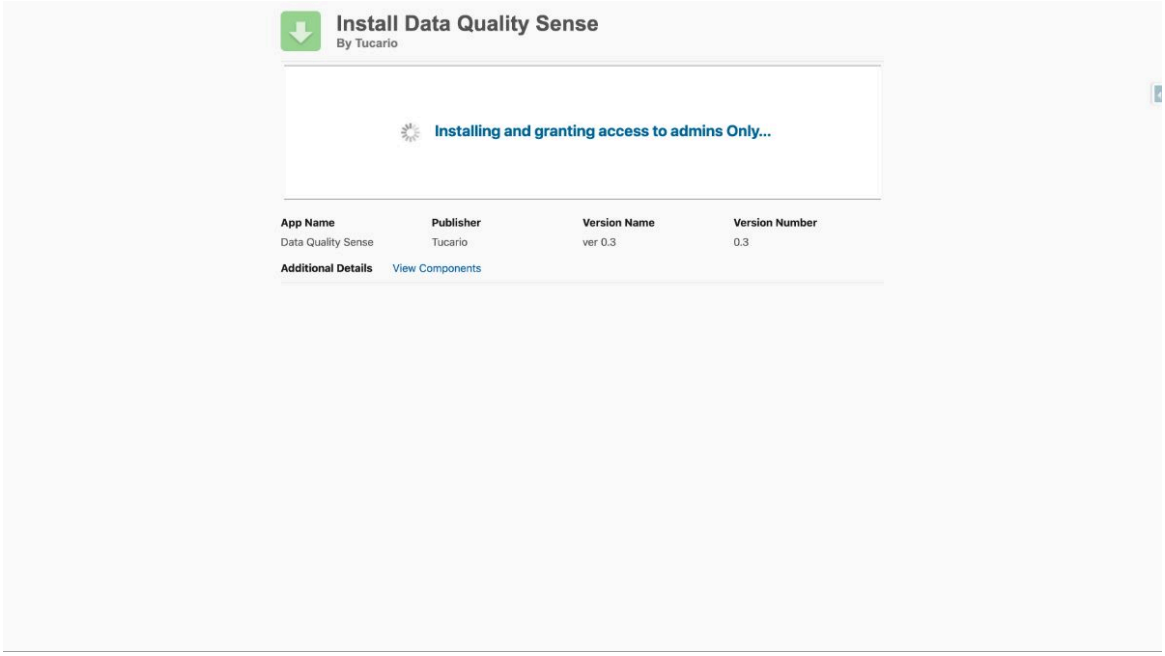

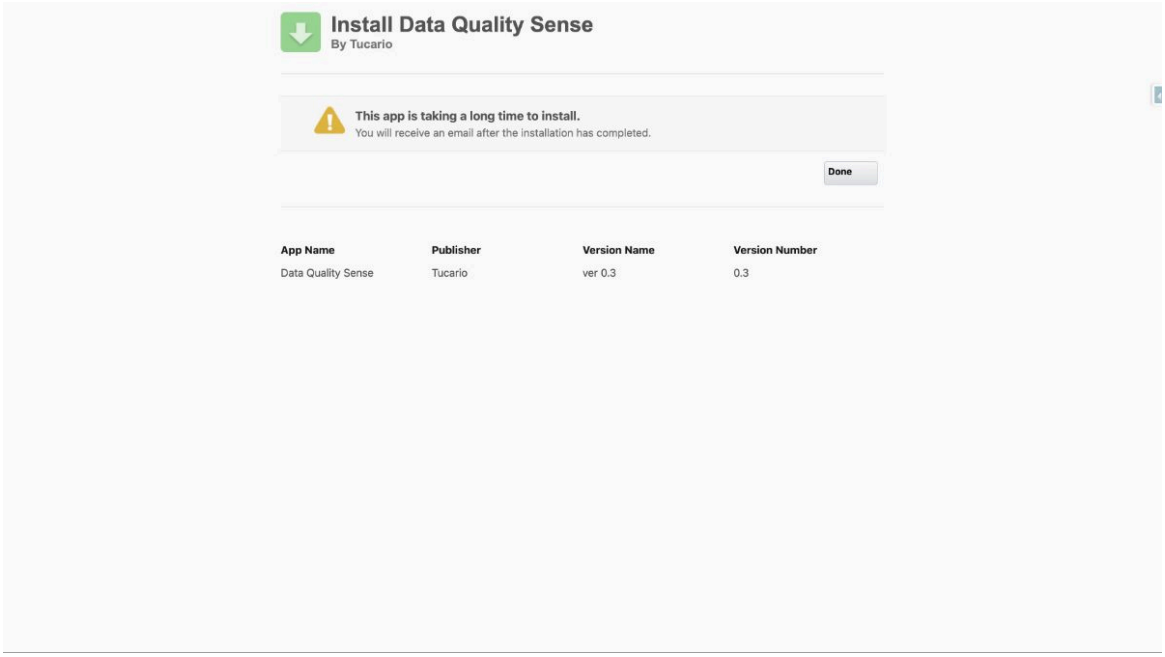

5. **Verify installation**

   Navigate to **Setup** → **Installed Packages** and confirm `Data Quality Sense` appears with namespace `dataqualitysens`.

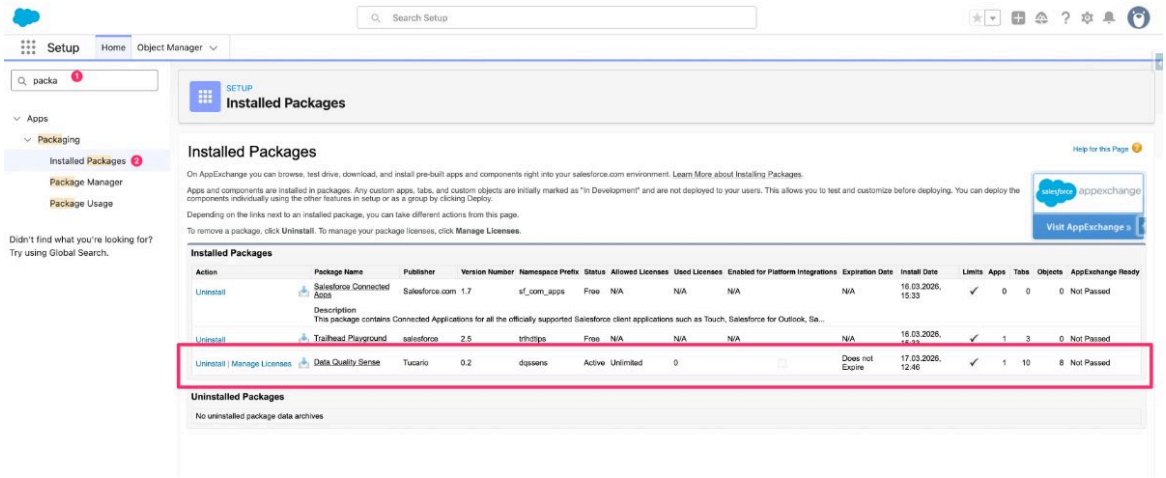

Tip

After installation, proceed to Permissions to configure user access.

Activate the App

Section titled “Activate the App”

After installation, Data Quality Sense is **not yet active**. When you open the app for the first time, you will see a lock screen indicating that activation is required.

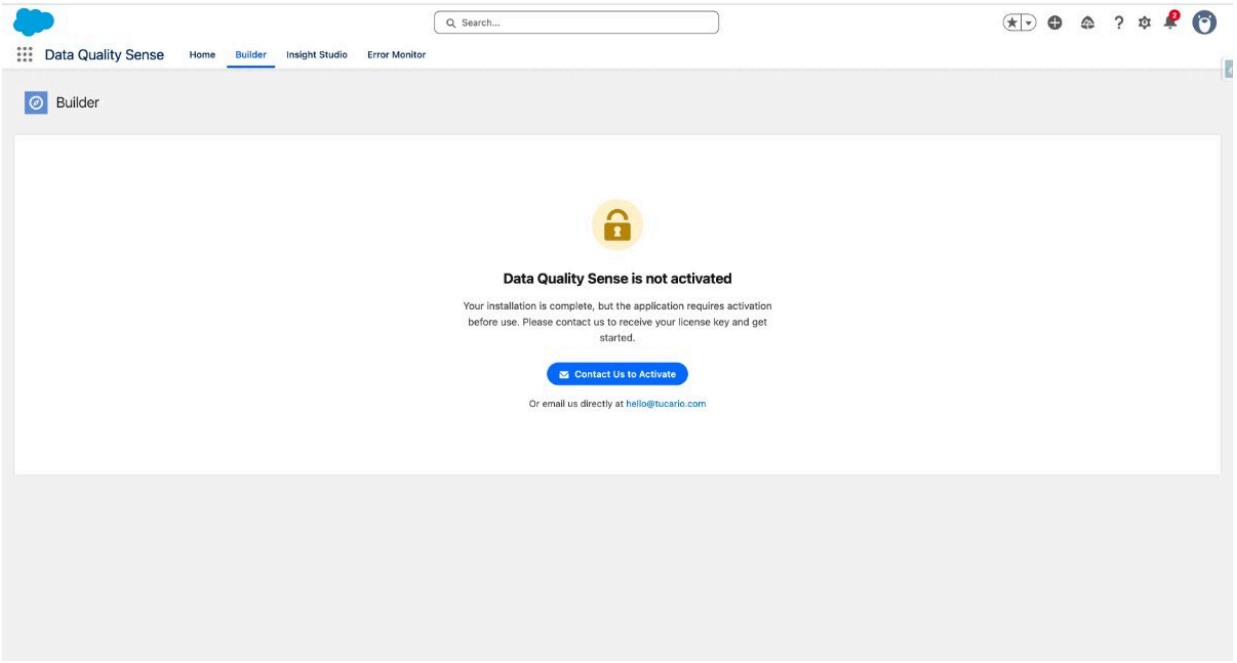

Click **Contact Us to Activate** to open a pre-filled email with your Org ID and details. You can also email us directly at hello@tucario.com.

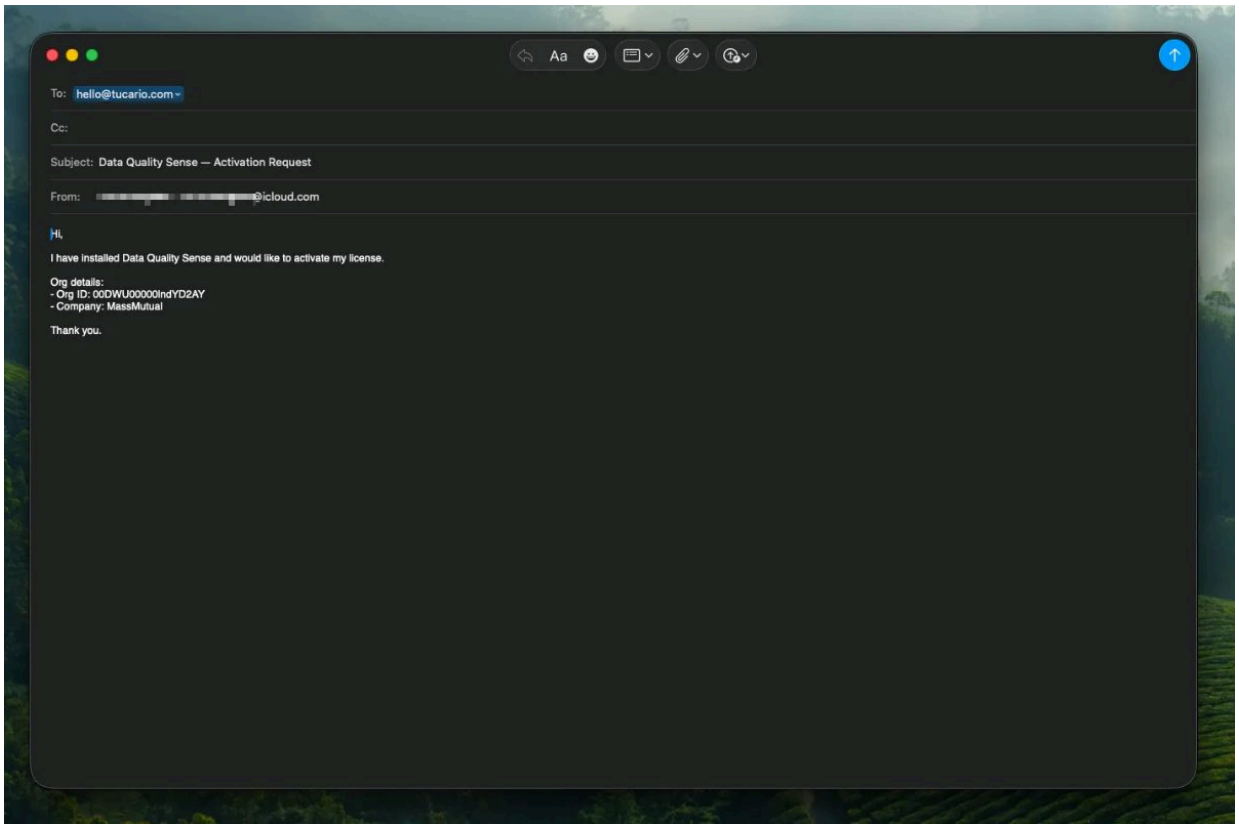

We will enable the app remotely — no additional installation steps are needed on your side. Once activated, the lock screen disappears and the app is ready to use with **20 scans** included.

ⓘNote

Activation is tied to your Salesforce org. If you install the package on a new org (e.g., a sandbox), it needs to be activated separately.

Post-Installation

🔗Section titled “Post-Installation”

The package installs all required components — custom objects, Lightning components, and permission sets. No additional configuration of these components is needed.

Updating the Package

🔗Section titled “Updating the Package”

To upgrade to a newer version, use the same installation link — Salesforce will detect the existing package and offer an upgrade path. The installer shows your currently installed version and the new version available. Choose the same installation scope as before (or change it if needed) and click **Upgrade**.

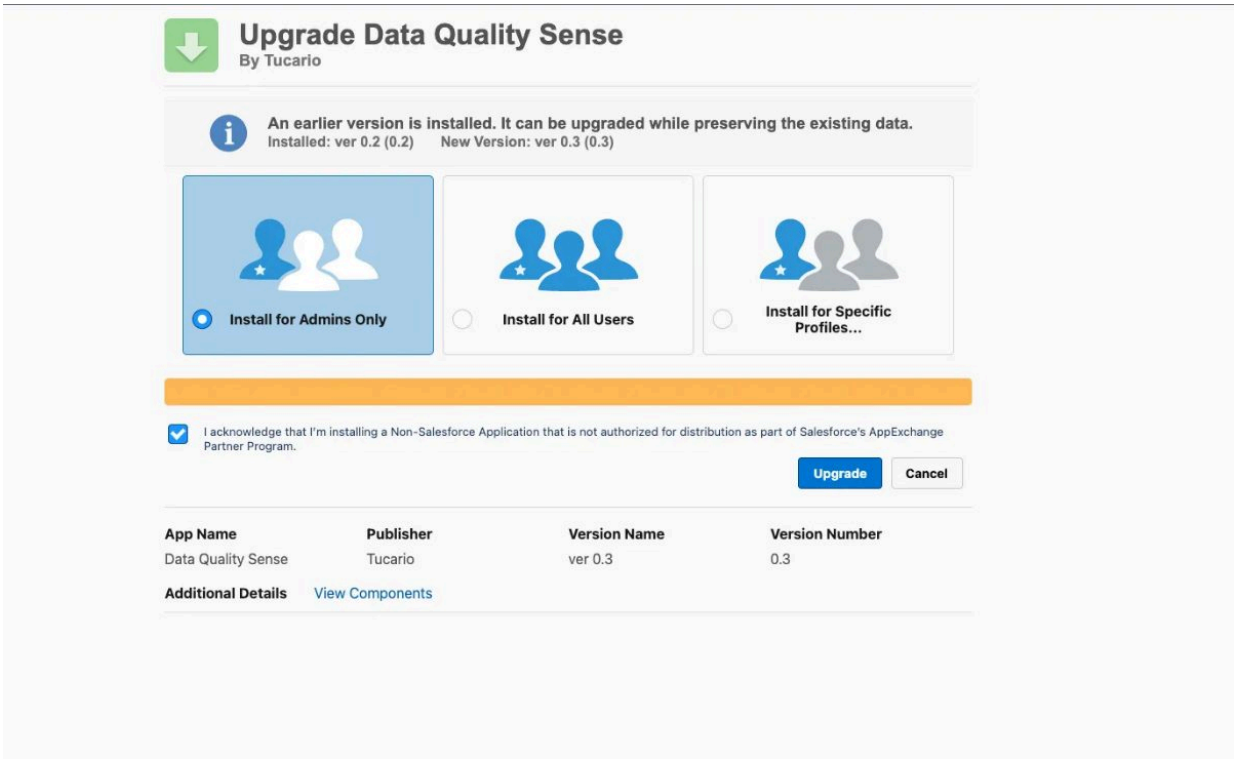

Your scan definitions, results, and configuration are preserved during upgrades — no data is lost.

Permissions

Permission Sets

Section titled “Permission Sets”

Data Quality Sense ships with permission sets that control access to different parts of the application.

DQS Admin

Section titled “DQS Admin”

Full access to all features:

- Create, edit, and delete scan definitions
- Configure capabilities and thresholds
- Run and schedule scans
- View all results in Insight Studio
- Access the Error Management Console
- Manage data retention settings

DQS User

Section titled “DQS User”

Read-only access to results:

- View scan results in Insight Studio
- Export data to CSV
- View scan schedules

Assigning Permission Sets

Section titled “Assigning Permission Sets”

1. Navigate to **Setup** → **Users** → **Permission Set Assignments** for the target user
2. Click **Edit Assignments**
3. In the **Available Permission Sets** list, find **DQS Admin** or **DQS User**
4. Move the desired permission set to **Enabled Permission Sets** using the arrow button
5. Click **Save**

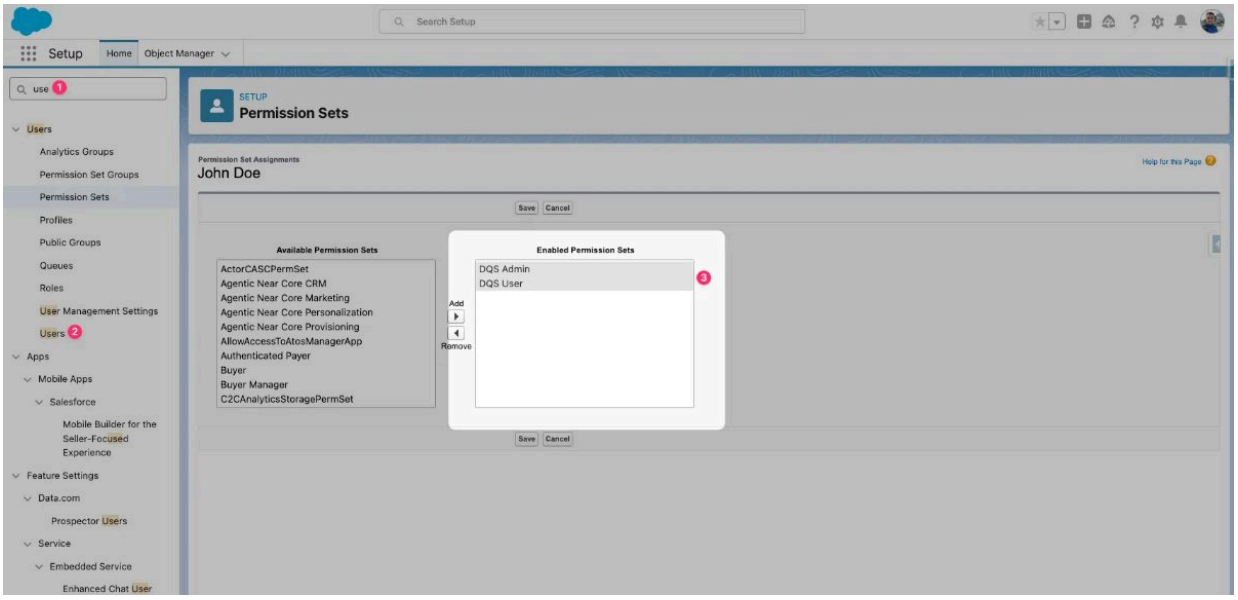

⚠Caution

Users without any DQS permission set assigned will see an activation gate when accessing the application. The Getting Started tab will appear automatically to guide administrators through the setup process.

Object Permissions

Section titled “Object Permissions”

The permission sets automatically grant access to DQS custom objects:

Object	Admin	User
DQS_Definition__c	Read/Create/Edit/Delete	Read
DQS_Definition_Detail__c	Read/Create/Edit/Delete	Read
DQS_Dimension_Result__c	Read	Read
DQS_Field_Result__c	Read	Read
DQS_Metric_Result__c	Read	Read
DQS_Batch_Schedule__c	Read/Create/Edit/Delete	Read

Field-Level Security

Section titled “Field-Level Security”

All fields on DQS objects are visible to both permission sets. Metric results, scores, and configuration fields are read-only for DQS User.

Quick Start

Your First Scan in 5 Minutes

Section titled “Your First Scan in 5 Minutes”

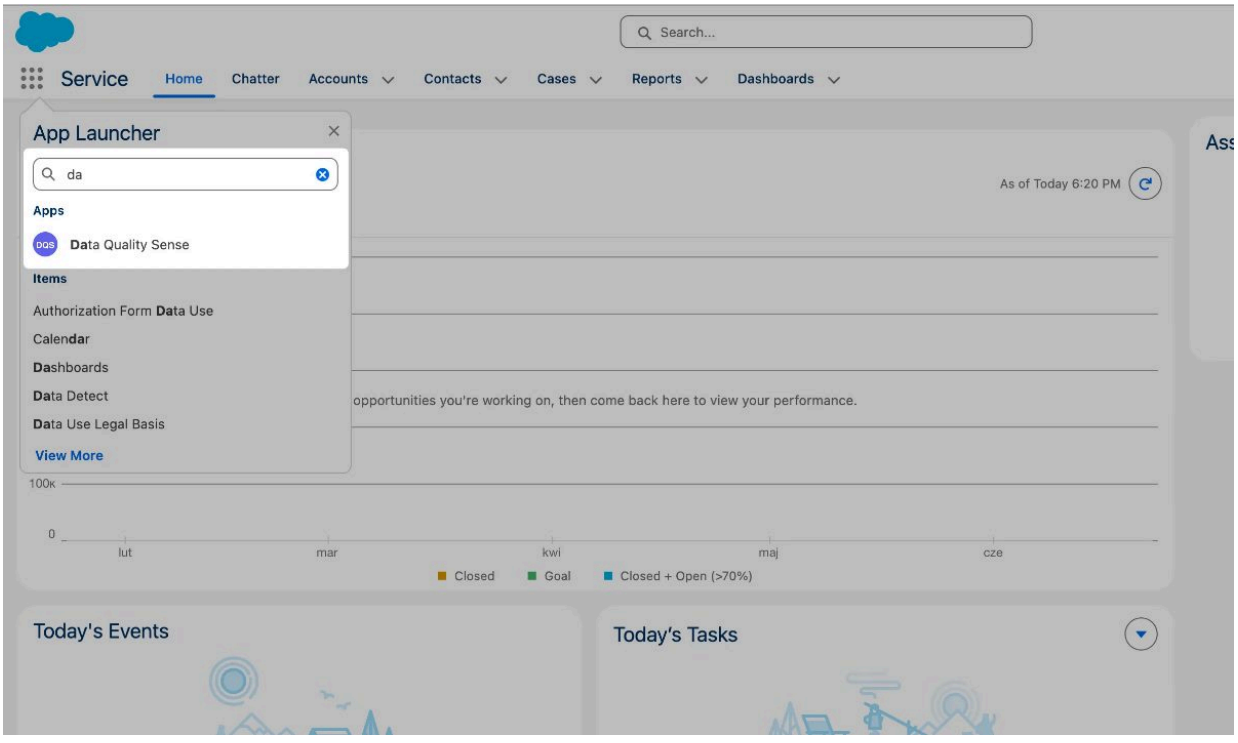

1. **Open DQS Builder**

   Navigate to the **Data Quality Sense** app in Salesforce. If permissions are configured correctly, the Builder tab loads automatically.

2. **Create a new definition**

   Click **New Definition**. Give it a name (e.g., “Account Quality Check”) and select the target object — for example, `Account`.

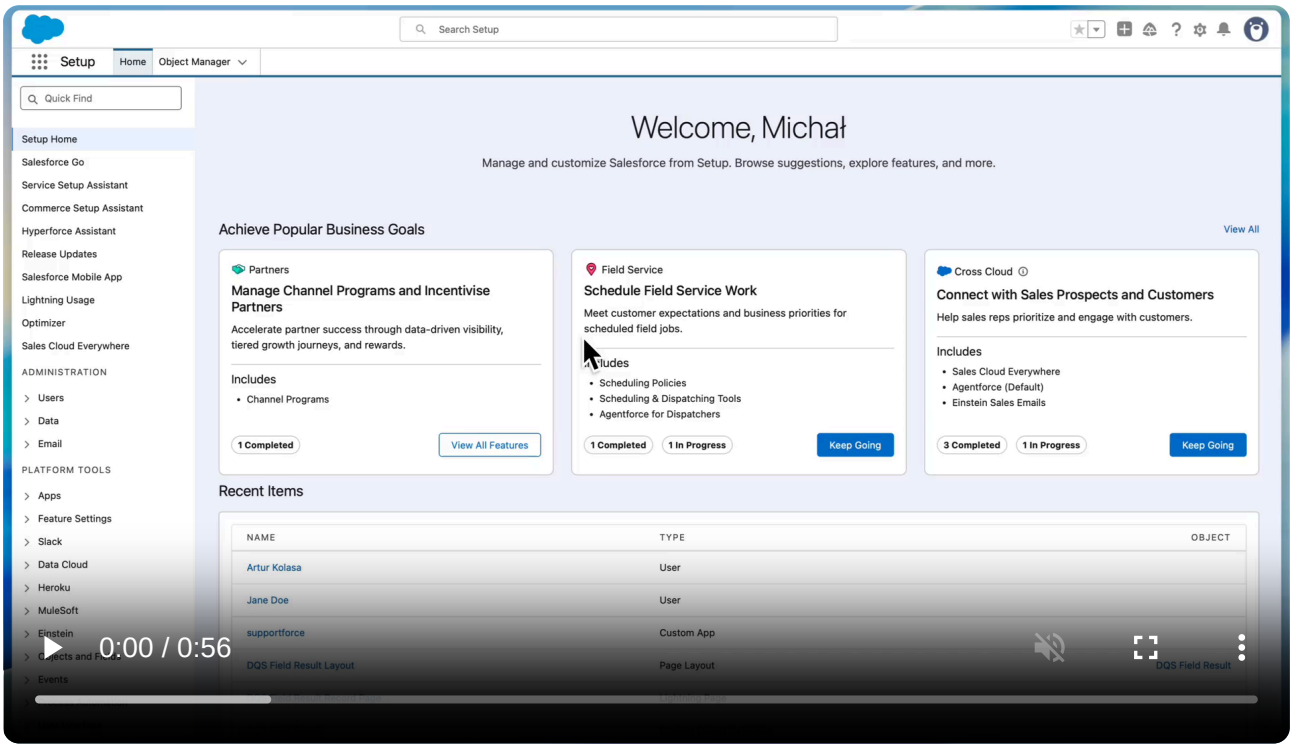

3. **Select fields**

   The field picker shows all available fields on the selected object. Pick the fields you want to monitor. You can use the search bar to filter, and sort by field type or label.

4. **Choose capabilities**

   Select which quality dimensions to evaluate:

   - **Completeness** — are the fields populated?
   - **Validity** — do values match expected formats?
   - **Uniqueness** — are there duplicate values?
   - **Timeliness** — is the data current and up-to-date?
   - **Consistency** — are related fields logically consistent?
   - **PII Detection** — does free text contain personally identifiable information?

   Start with Completeness for the quickest setup.

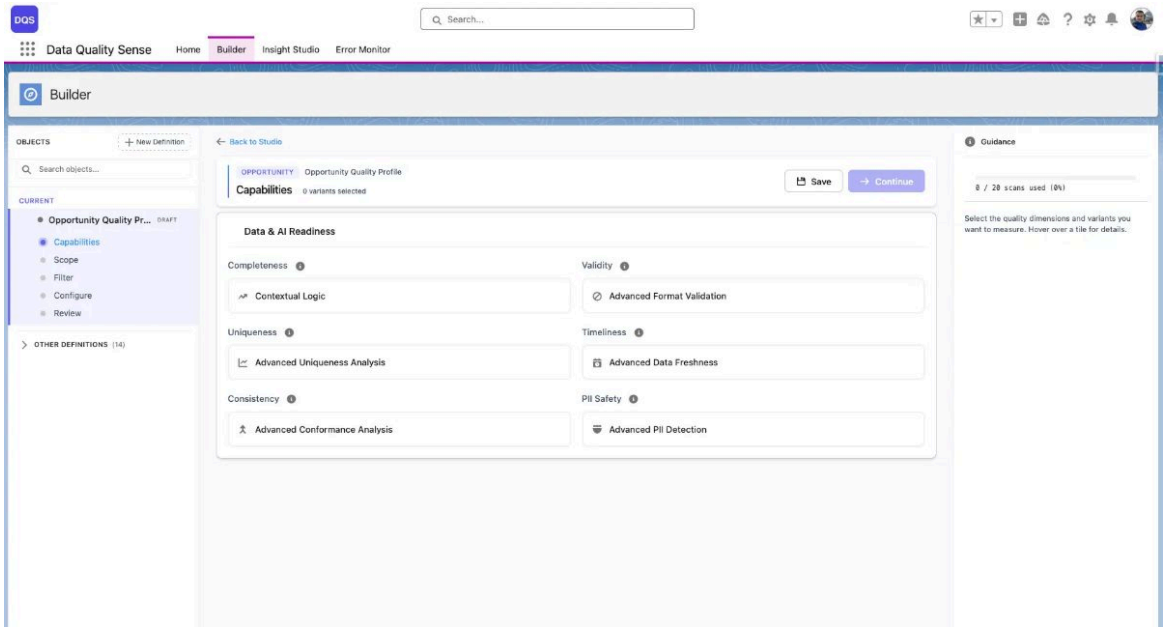

5. **Review and activate**

   Review your configuration on the summary screen. Click **Activate** to move the definition from Draft to Active status.

6. **Run the scan**

   Go to **Insight Studio** and trigger a scan manually, or set up a schedule.

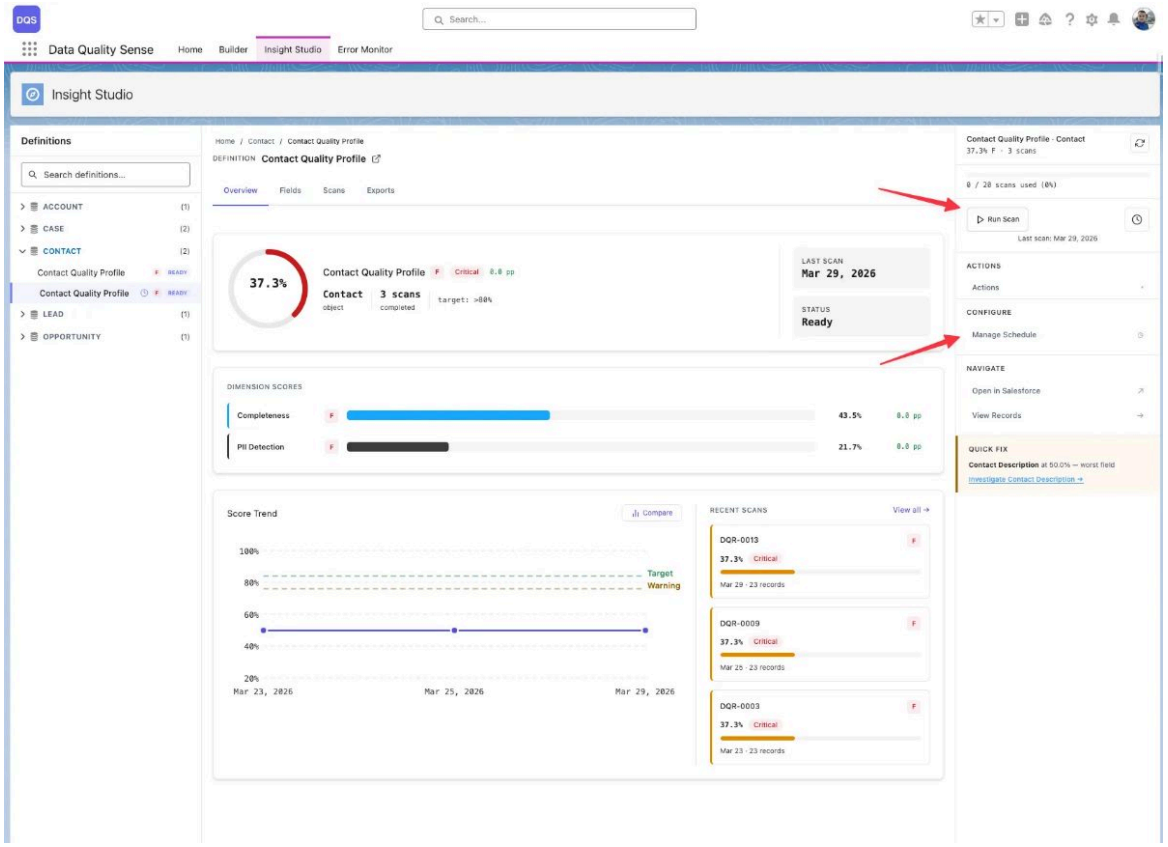

7. **View results**

After the scan completes, Insight Studio shows your data quality scores with drill-down to field-level details.

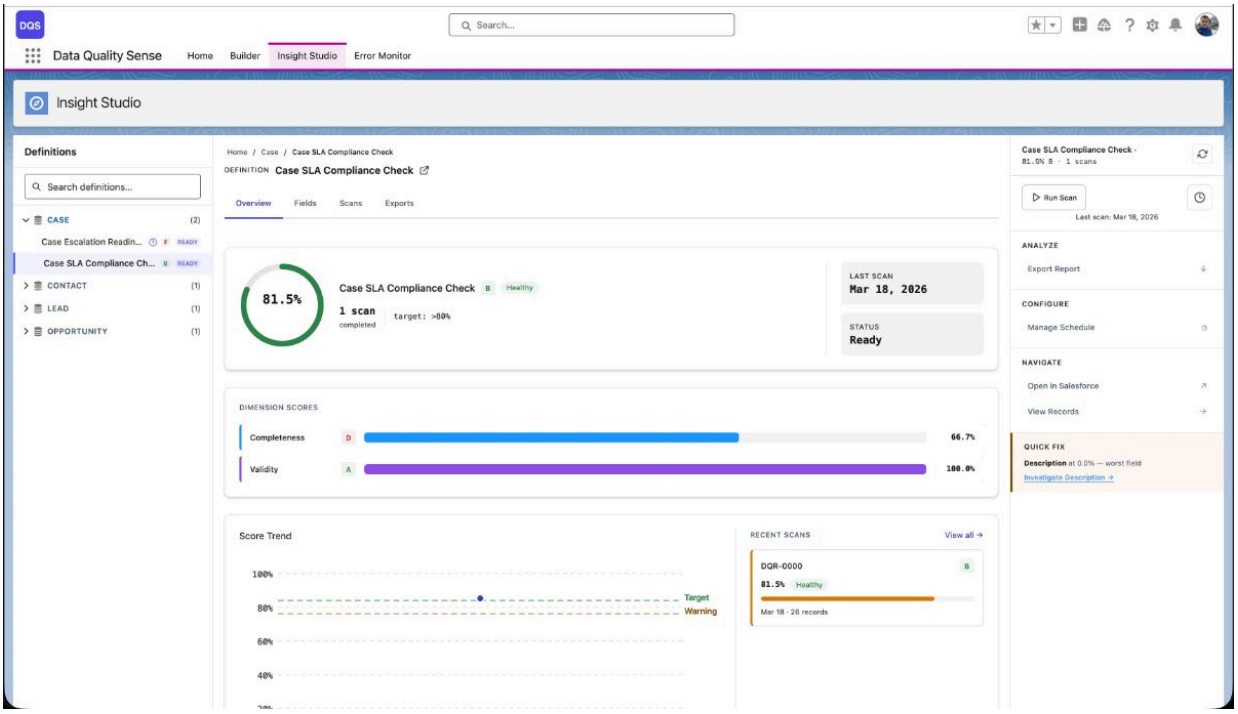

🚀Tip

You can always go back and add more capabilities or fields to your definition later. The Builder supports editing Active definitions.

What’s Next?

🔗Section titled “What’s Next?”

- Learn about all 7 quality capabilities
- Explore Insight Studio dashboards
- Set up automated scan schedules

Builder

Builder Overview

What is the Builder?

Section titled “What is the Builder?”

The DQS Builder is a **multi-step configuration wizard** that lets you define exactly how your data quality should be measured. It guides you through selecting objects, picking fields, choosing quality dimensions, and setting thresholds.

Builder Workspace

Section titled “Builder Workspace”

The Builder uses a **3-zone layout**:

- **Sidebar** (orange) — Navigation tree showing your definitions and their configuration stages (Capabilities, Scope, Fields, Configure, Review)
- **Stage** (green) — Main content area where you configure each step of the definition
- **Guidance Panel** (red) — Contextual help and guidance that adapts to your current step

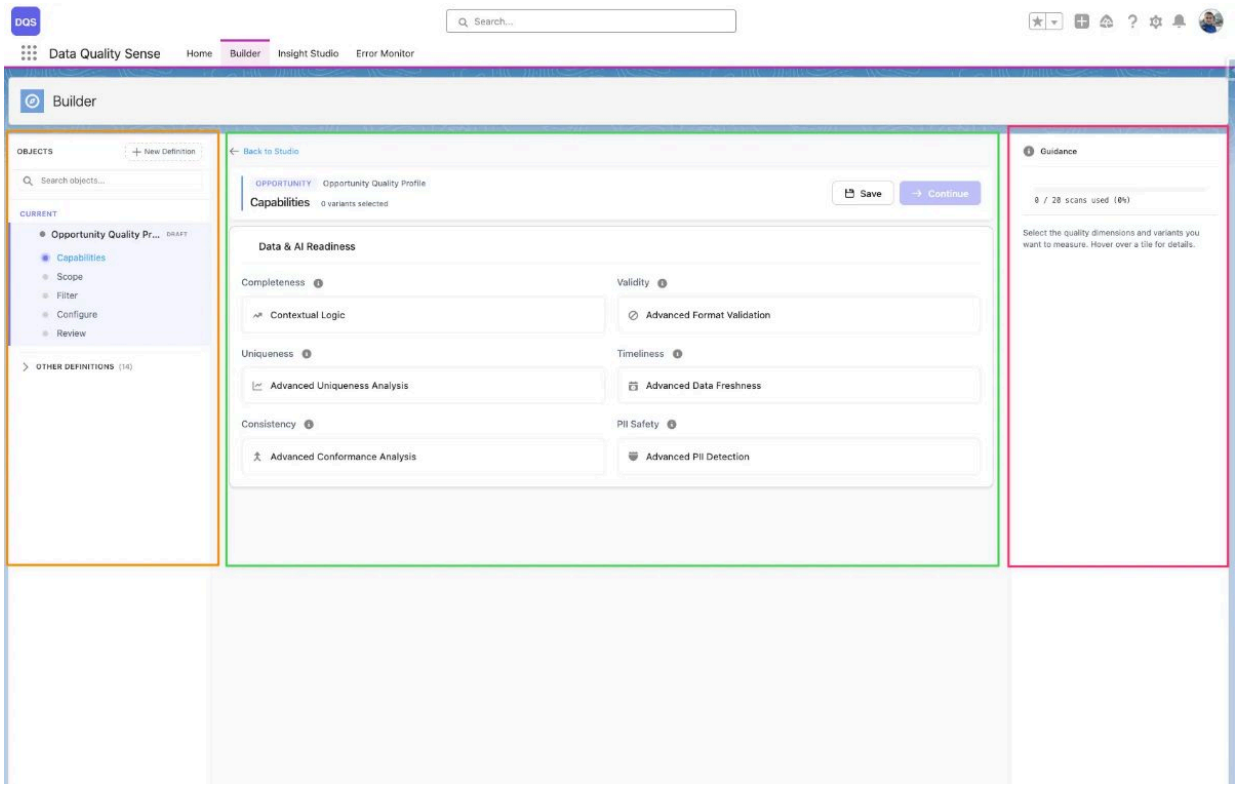

Key Concepts

Section titled “Key Concepts”

Definition

A scan configuration that specifies which object, fields, and quality capabilities to evaluate. Each definition produces a set of results when scanned.

Capability

A quality dimension like Completeness or Validity. Each capability can be configured globally or overridden per field.

Definition Detail

A record linking a specific field to a definition, storing per-field configuration overrides for each capability.

Lifecycle

Definitions progress through stages: Draft → Ready → Active → Obsolete. Only Active definitions can be scanned.

Builder Stages

Section titled “Builder Stages”

1. **Getting Started** — Select the target Salesforce object
2. **Field Selection** — Pick which fields to include in the scan
3. **Capability Configuration** — Enable and configure quality dimensions

4. **Summary** — Review the complete definition before activation

Each stage validates your input before allowing you to proceed to the next step.

Learn More

Section titled “Learn More”

- Creating a Definition
- Field Selection
- Configuring Capabilities
- Definition Lifecycle

Creating a Definition

What is a Definition?

Section titled “What is a Definition?”

A **definition** is the core configuration unit in DQS. It specifies:

- Which Salesforce object to scan
- Which fields to evaluate
- Which quality dimensions (capabilities) to measure
- Optional per-field threshold overrides

Creating a New Definition

Section titled “Creating a New Definition”

1. **Click “New Definition”**

   In the Builder tab, click the **New Definition** button in the top-left corner. A creation dialog opens with object search, recent definitions, and suggested objects.

2. **Select the target object**

   Pick the SObject you want to scan. You can search all objects in your org, choose from **Recent Definitions** (objects you’ve already built definitions for), or select from the **Suggested** list (Account, Opportunity, Contact, Case, Lead). The dialog shows how many records and existing definitions each object has.

3. **Enter a definition name**

   Once you select an object, the right-hand panel shows a name field with auto-generated **suggestions** based on the object (e.g., “Opportunity Quality Profile”, “Opportunity Completeness Audit”). Pick a suggestion or type your own name.

4. **Click ”+ Create”**

   The definition is created in **Draft** status and opens in the Builder wizard.

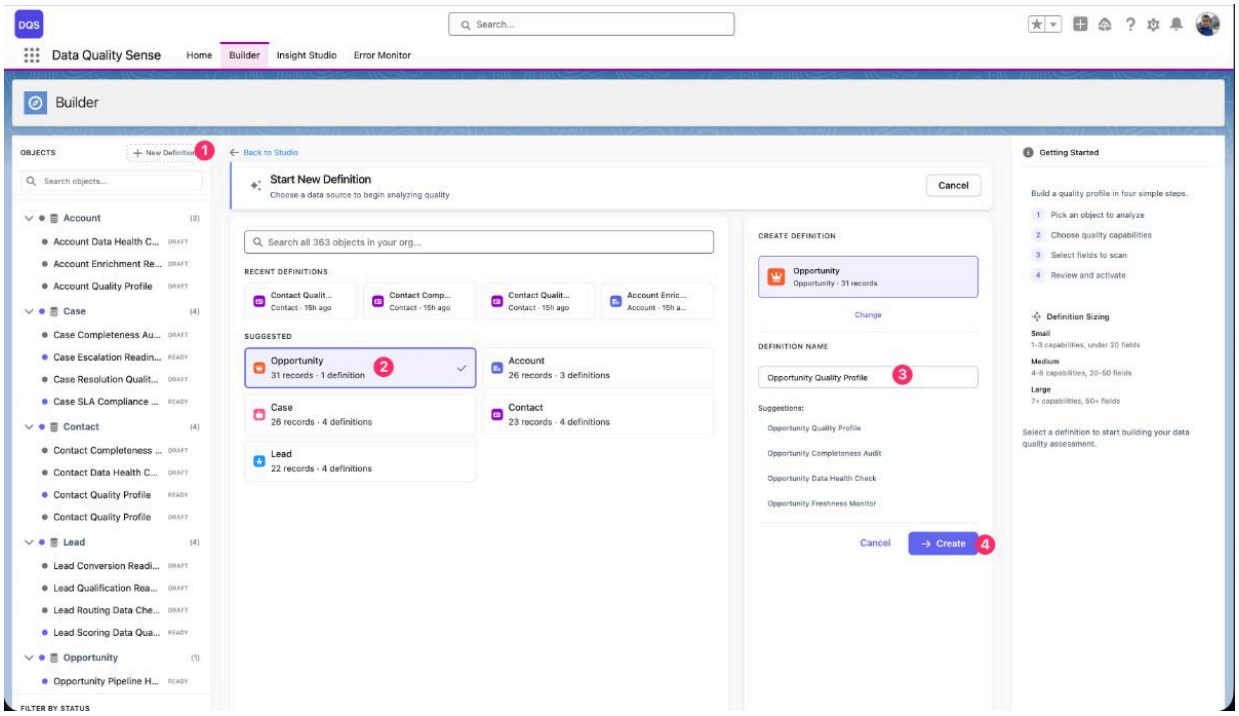

ⓘNote

Technical objects (those with API names starting with common internal prefixes) are filtered out by default to keep the object list manageable.

Definition Settings

🔗Section titled “Definition Settings”

After creation, you can configure:

Setting	Description
Name	Display name shown in lists and Insight Studio
Object	The target SObject (cannot be changed after creation)
Description	Optional notes about the scan’s purpose
Status	Current lifecycle stage (Draft, Ready, Active, Obsolete)

Editing Existing Definitions

Section titled “Editing Existing Definitions”

Open any definition from the **Home** tab’s recent activity table, or from the Builder’s sidebar navigation tree. Definitions in **Draft** status can be edited — click **Continue Building** to resume configuration.

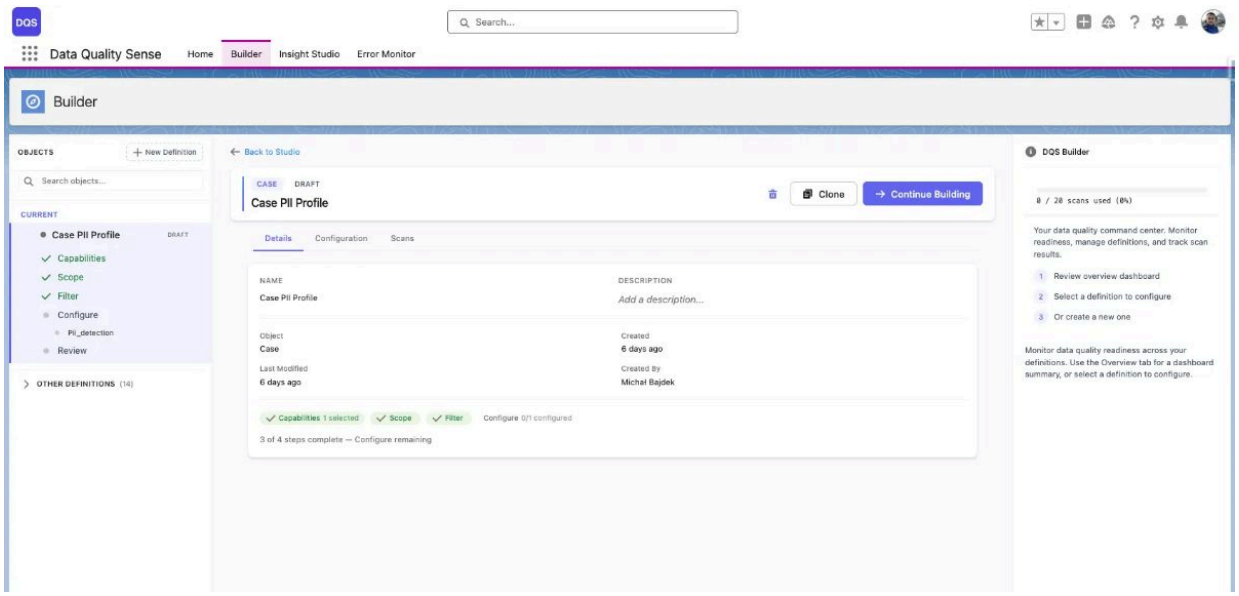

Deleting Definitions

Section titled “Deleting Definitions”

The delete button (trash icon) is only available when **both** conditions are met:

- The definition is in **Draft** status
- You have the **DQS Admin** permission set assigned

Active or completed definitions cannot be deleted — retire them to Obsolete status instead.

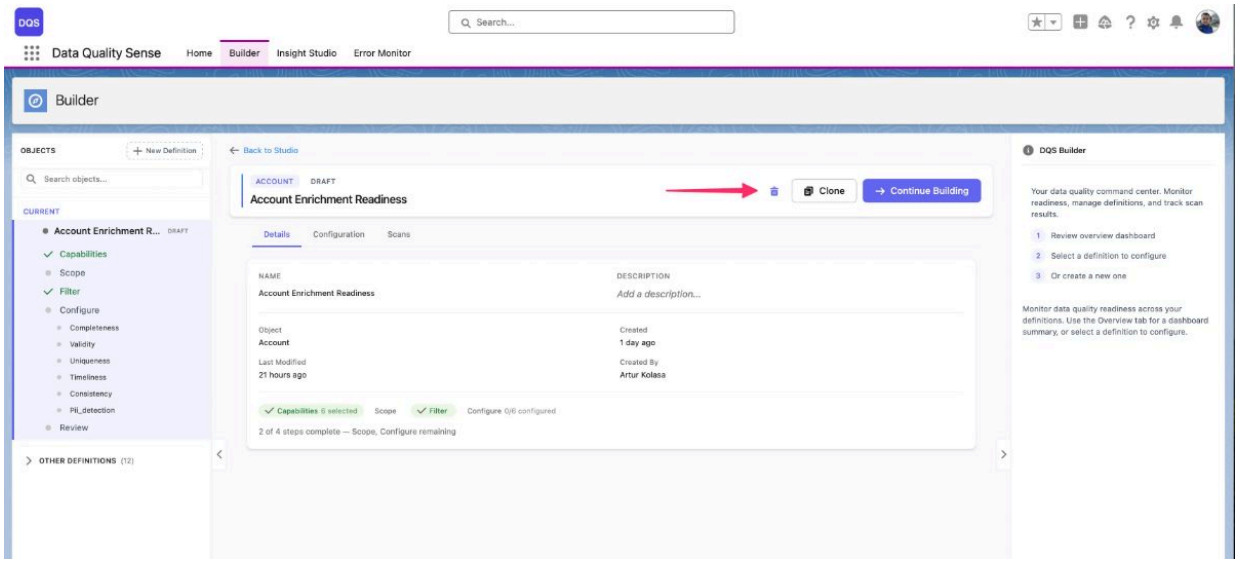

Deleting a definition removes all associated definition details, but preserves historical scan results for audit purposes.

Field Selection

Field Picker

Section titled “Field Picker”

The field picker is the second stage of the Builder wizard. It displays all available fields on your selected object in a paginated table.

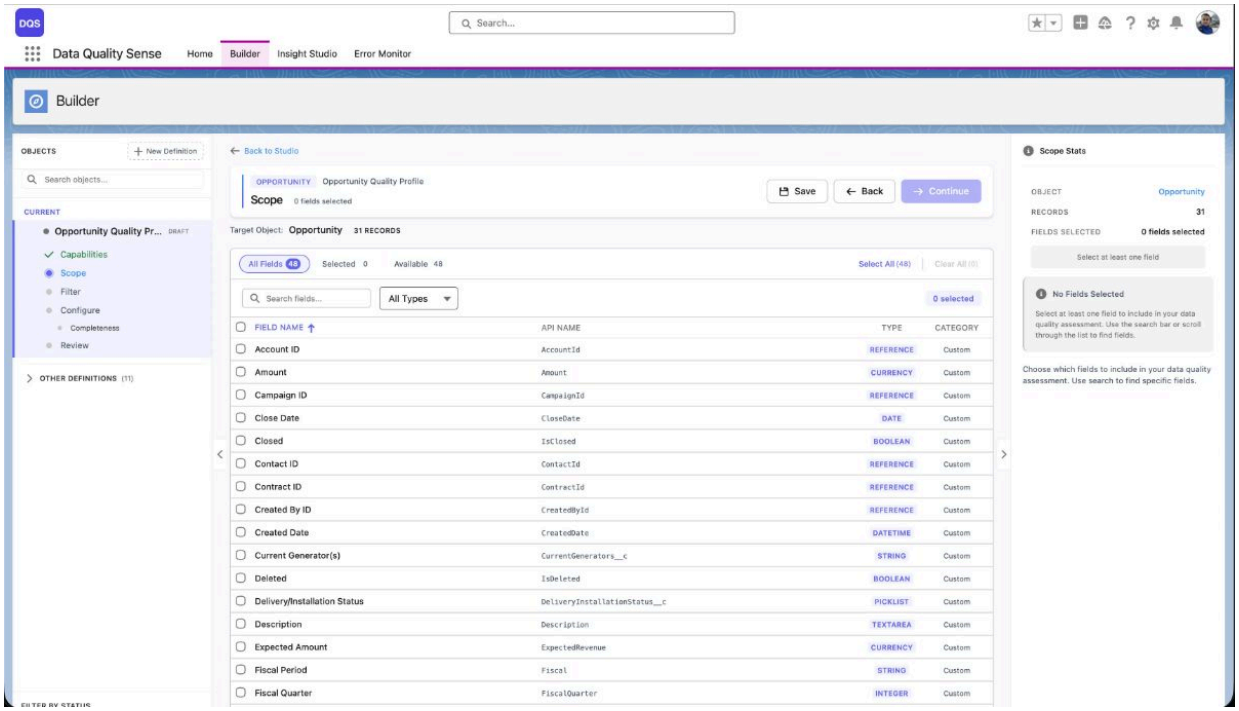

Features

Section titled “Features”

Search and Filter

Section titled “Search and Filter”

- **Search** — Filter fields by label or API name
- **Field Type Filter** — Show only specific field types (Text, Number, Picklist, etc.)
- **Sortable Columns** — Sort by field label, API name, or field type

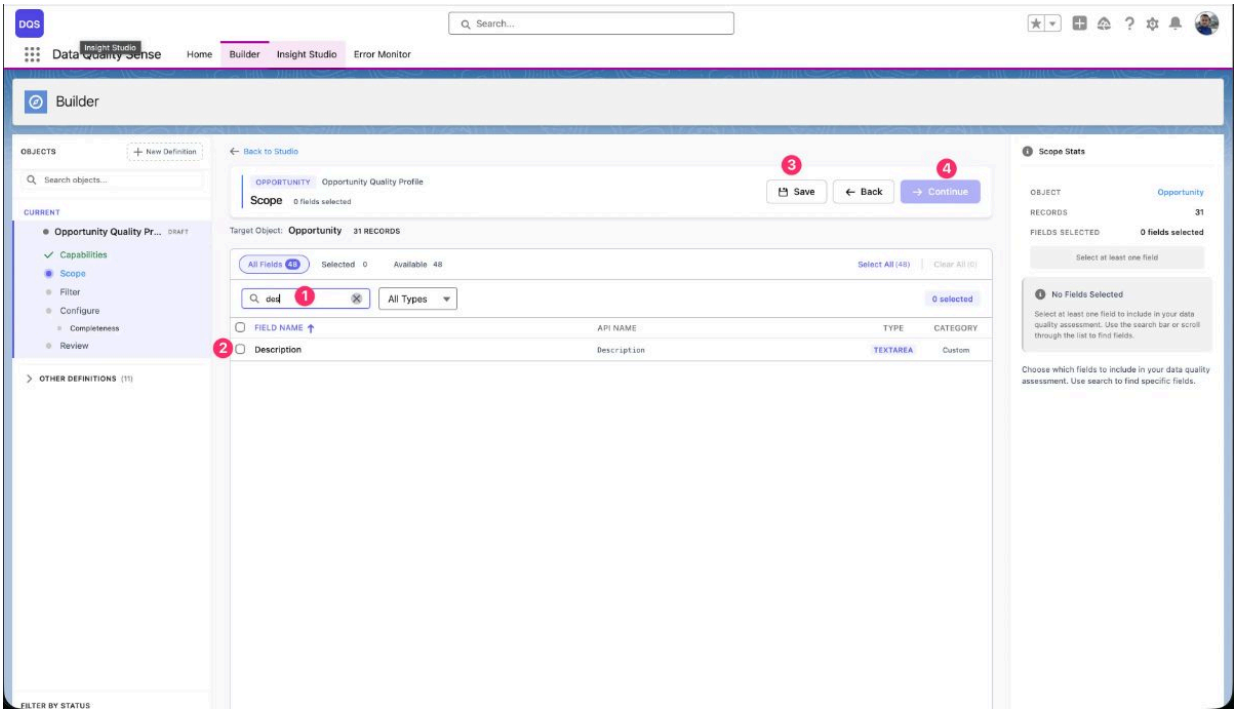

Field Information

Section titled “Field Information”

For each field, the picker displays:

Column	Description
Label	The field’s display name
API Name	The developer name (e.g., `BillingCity`)
Type	Field data type (Text, Number, Date, Picklist, etc.)
Required	Whether the field is required on the layout

Bulk Selection

Section titled “Bulk Selection”

Use the **Select All** link to add all available fields to your scope at once — the link shows the total field count (e.g., “Select All (71)”). To remove all selections, click **Clear All**. The header bar displays live counters for **Selected** and **Available** fields so you always know how many fields are in scope.

Selections persist across pages — you can select fields on page 1, navigate to page 2, and your selections are preserved. The **Scope Stats** panel on the right summarizes your current selection.

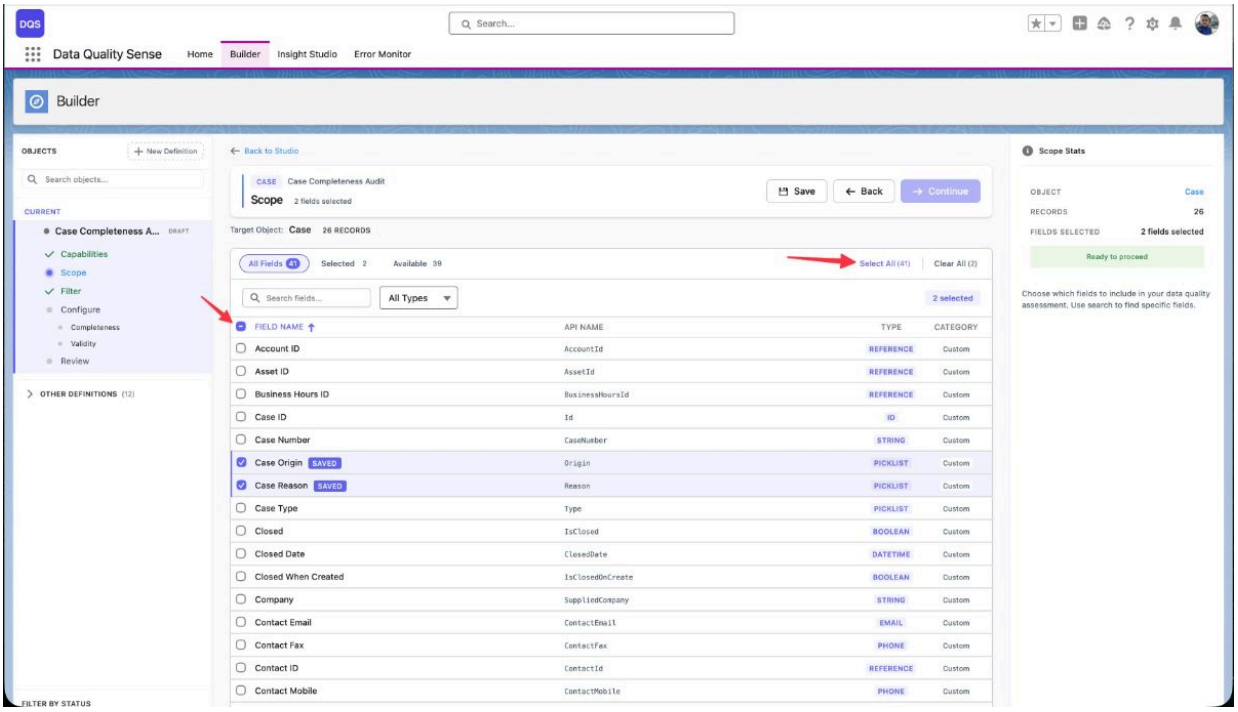

Field Considerations

Section titled “Field Considerations”

Tip

Not all field types are applicable to all capabilities. For example, Uniqueness may not be meaningful for Boolean fields. DQS automatically handles non-applicable combinations.

Recommended Fields

Section titled “Recommended Fields”

For a comprehensive scan, include:

- **Key identifier fields** — Name, Email, Phone
- **Address fields** — for Completeness and Validity
- **Date fields** — for Timeliness
- **Picklist fields** — for Validity and Consistency
- **Free text fields** — for Completeness and PII Detection

Fields to Avoid

Section titled “Fields to Avoid”

- System-generated fields (Id, CreatedDate, LastModifiedDate) — these are always populated
- Formula fields — their values are computed, not entered by users
- Encrypted fields — may not be readable by the scan engine

Scope View

Section titled “Scope View”

The Scope view is the main workspace for selecting and reviewing fields:

1. **Scope step** — Click the **Scope** step in the Builder wizard sidebar to load all fields for the target object. This is where you define which fields will be included in your data quality scan.
2. **All Fields** — The default view showing every field on the object. Use this tab to browse the full list, check or uncheck fields, and see each field’s label, API name, type, and whether it’s a standard or custom field. Fields you select are highlighted in the list.
3. **Selected filter** — Switch to this view to see only the fields you’ve already added to scope. The counter shows how many fields are selected (e.g., “Selected: 2”). Use it to quickly review your choices and remove any fields you no longer need.
4. **Available filter** — Shows only the fields that are not yet in your scope. The counter displays how many remain (e.g., “Available: 16”). Useful when you want to browse what’s left to add without scrolling past already-selected fields.
5. **Type filter** — The **All Types** dropdown lets you narrow the field list to a specific data type — Text, Number, Date, Picklist, Boolean, Lookup, and more. Combine it with the other filters to quickly find the exact fields you need (e.g., show only available Picklist fields).

The **Scope Stats** panel on the right summarizes your selection with a field count and shows a **Ready to proceed** indicator when at least one field is selected.

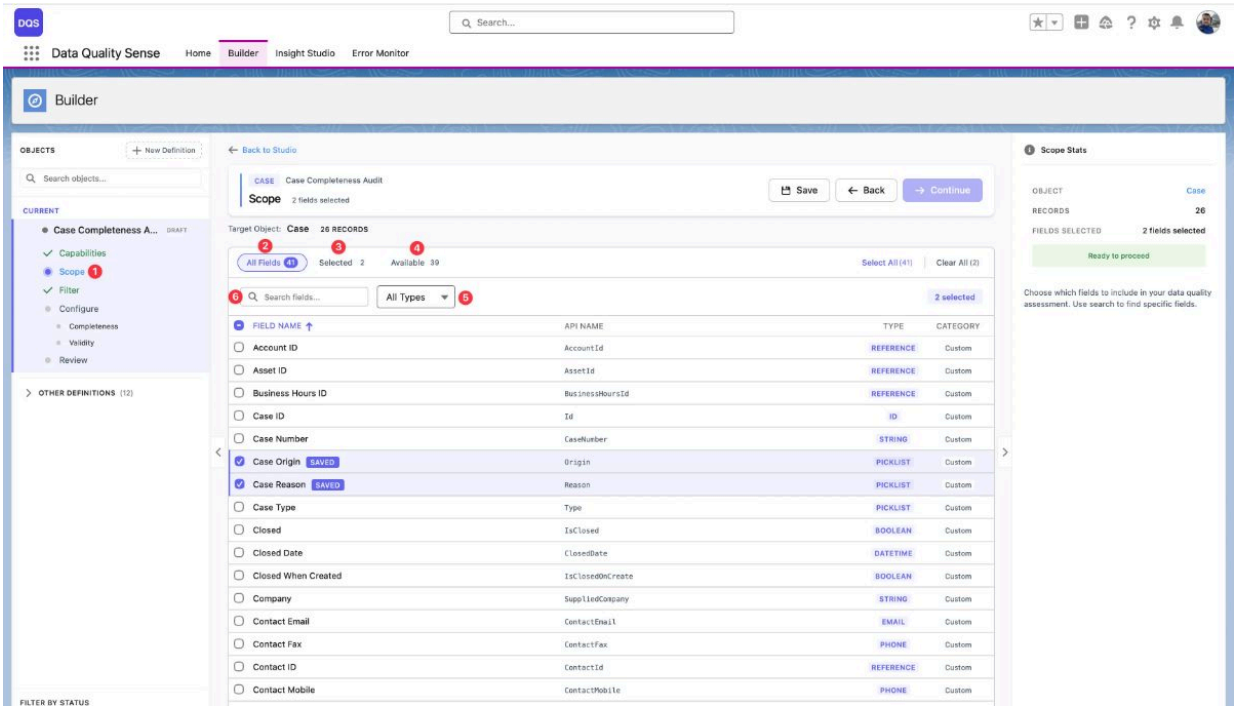

Configuring Capabilities

Capability Configuration

Section titled “Capability Configuration”

The third stage of the Builder wizard lets you enable quality dimensions and configure their settings. Each capability evaluates a different aspect of your data quality.

Available Capabilities

Section titled “Available Capabilities”

Capability	What It Measures
Completeness	Are fields populated?
Validity	Do values match expected formats?
Uniqueness	Are there duplicate values?
Timeliness	Is data up to date?
Consistency	Are related fields logically consistent?
PII Detection	Is personal data properly handled?

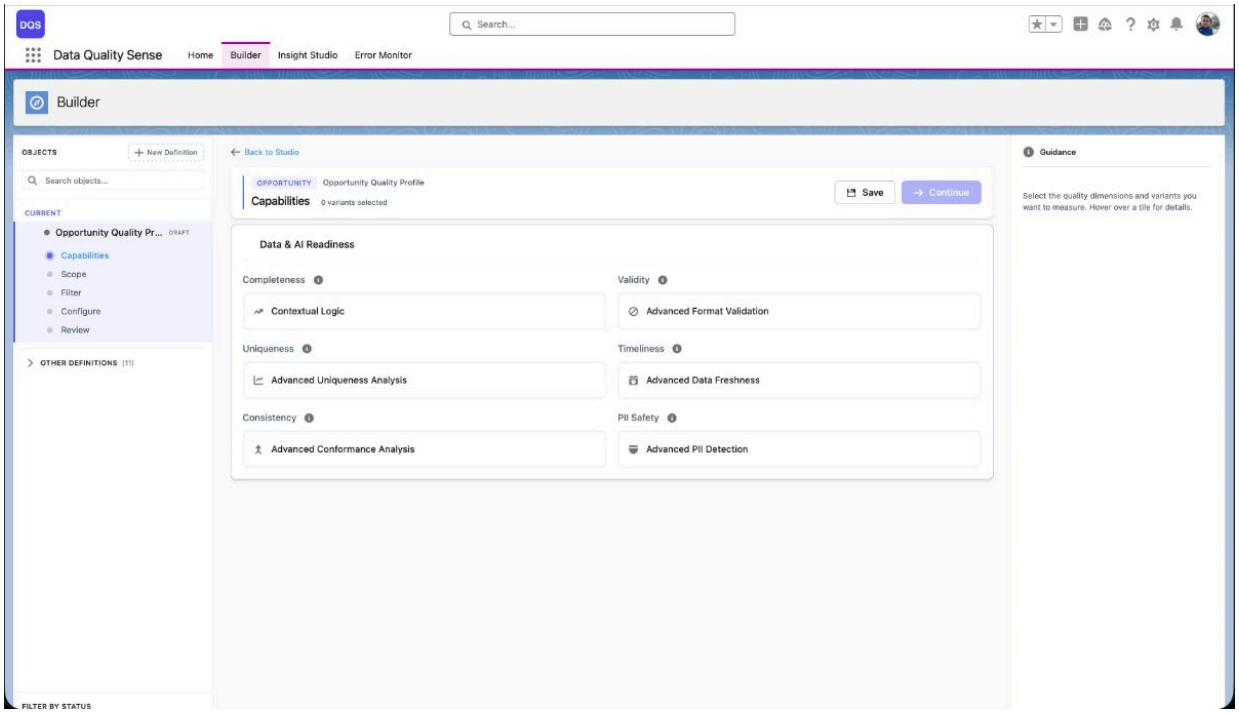

Enabling a Capability

Section titled “Enabling a Capability”

Select a capability from the list to enable it. Each capability has:

- **Global configuration** — Default settings applied to all fields
- **Per-field overrides** — Custom settings for individual fields that differ from the global defaults

Field Compatibility

Section titled “Field Compatibility”

Some capabilities require specific field types. For example, Timeliness only works with Date and DateTime fields. If none of the fields in your definition match, the capability displays a “Not Applicable” message and cannot be configured until you go back to scope and add compatible fields.

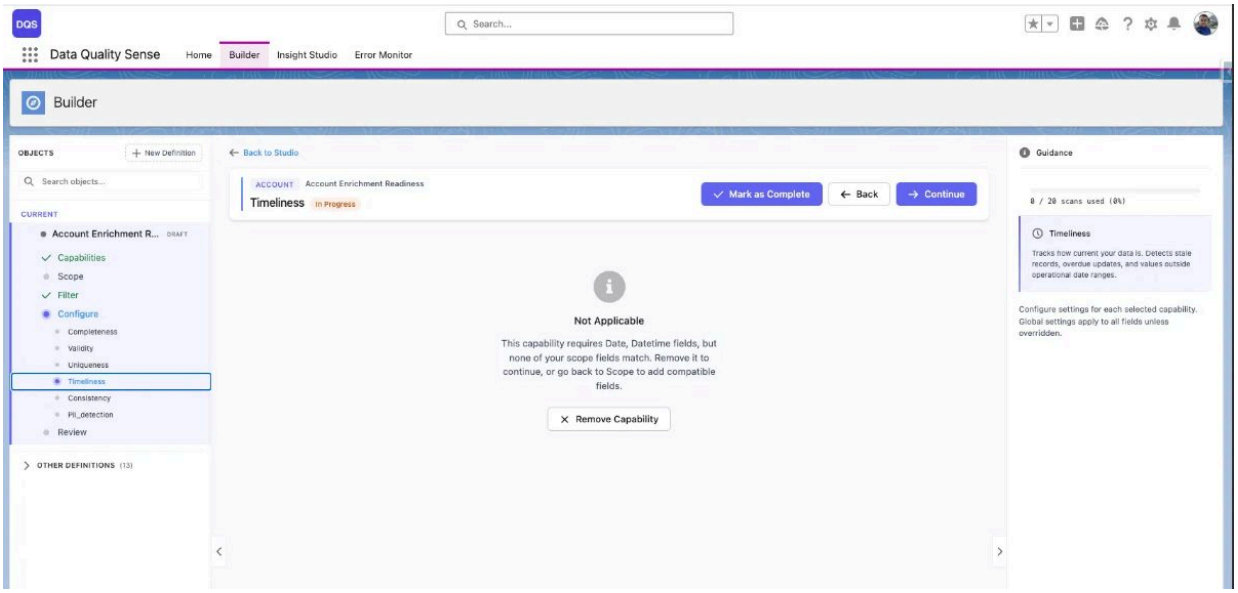

Global vs. Per-Field Configuration

Section titled “Global vs. Per-Field Configuration”

Most capabilities allow you to set global thresholds and then override them for specific fields. For example:

- **Completeness**: Set a global “expected fill rate” of 90%, but require 100% for Email fields
- **Validity**: Define global format rules, but set custom regex patterns for specific text fields
- **Timeliness**: Set a global freshness window of 30 days, but use 7 days for critical date fields

The **Defaults** section at the top controls global settings for the capability (e.g., Blank Handling and Placeholder Detection for Completeness). These apply to all fields unless overridden. The **Field Overrides** table below lists each field in scope with its current status — “Default” means the field uses the global settings.

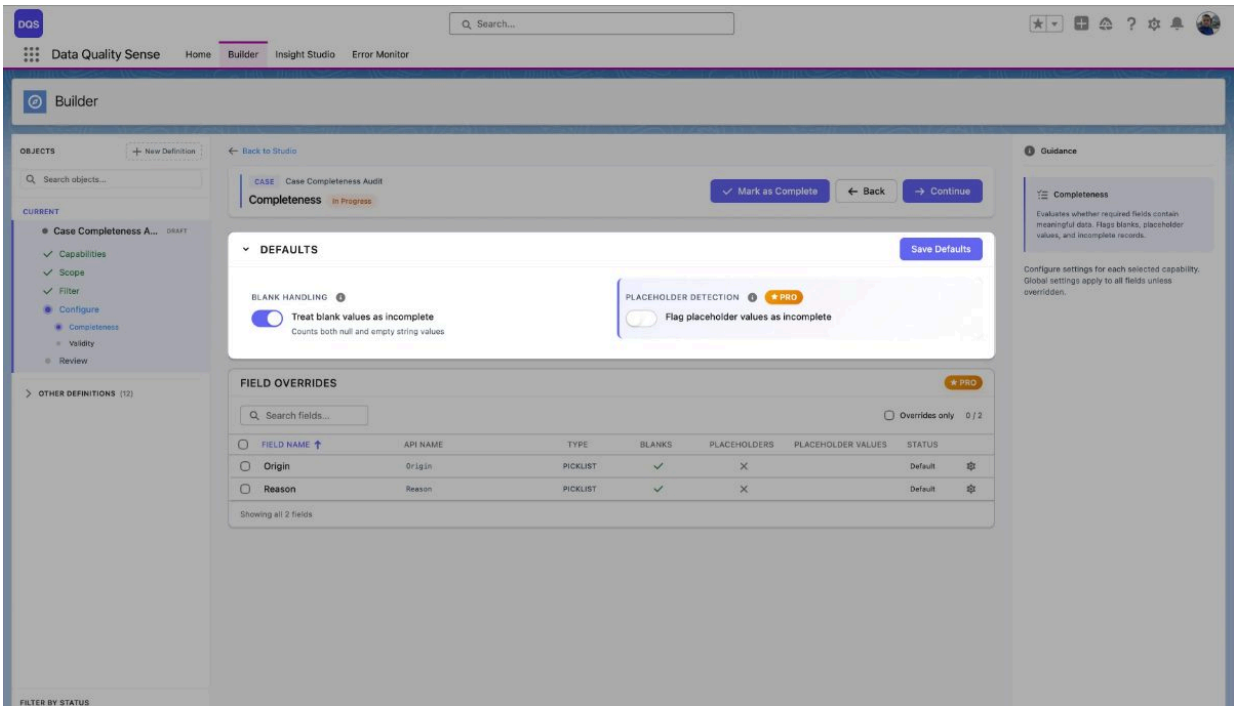

To override a specific field, click on it in the Field Overrides table. Its status changes to indicate a custom configuration. The red arrow below points to the Reason field after applying a per-field override.

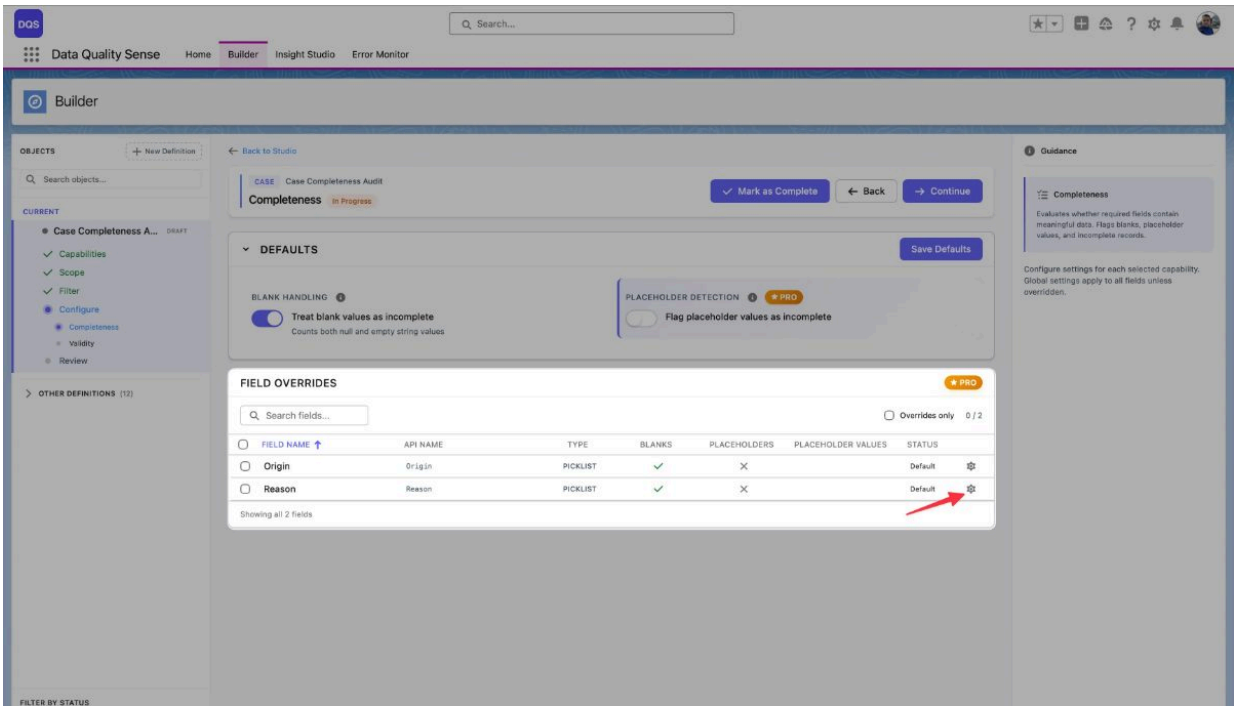

Configuration Workflow

[Section titled “Configuration Workflow”]

1. **Enable the capability** — Toggle it on from the capability list
2. **Set global settings** — Configure defaults that apply to all selected fields
3. **Add field overrides** (optional) — Click on individual fields to customize their settings
4. **Review** — The summary shows which fields have overrides

Tip

You can use **bulk configuration** for capabilities like Completeness to quickly set the same override across multiple fields at once.

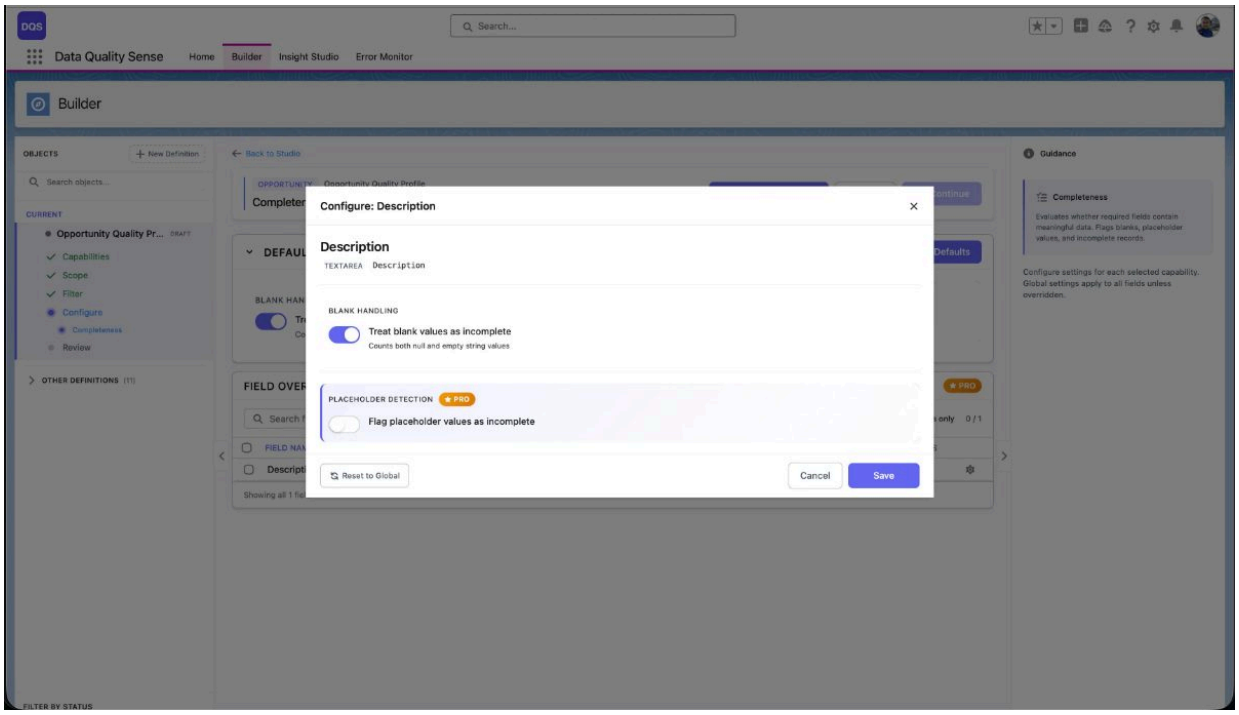

Removing a Capability

Section titled “Removing a Capability”

Click the remove button on any enabled capability to disable it. This clears all global and per-field configurations for that capability. The action requires confirmation.

Definition Lifecycle

Lifecycle Stages

Section titled “Lifecycle Stages”

Every scan definition progresses through a series of stages:

```
Draft → Ready → Active → Obsolete
```

Draft

Section titled “Draft”

The initial state when a definition is created. In Draft status:

- The definition can be freely edited
- Fields and capabilities can be added or removed
- No scans can be executed
- The definition is not visible in Insight Studio

Ready

Section titled “Ready”

A transitional state indicating the definition has been reviewed. In Ready status:

- All required configuration is complete
- The definition has passed validation
- It can be activated

Active

Section titled “Active”

The operational state. In Active status:

- Scans can be executed against the definition
- Results appear in Insight Studio
- The definition can be scheduled for recurring scans
- Editing is still possible — changes take effect on the next scan

Obsolete

Section titled “Obsolete”

The retired state. In Obsolete status:

- No new scans can be triggered
- Existing results remain visible in Insight Studio for historical analysis
- The definition can be reactivated if needed

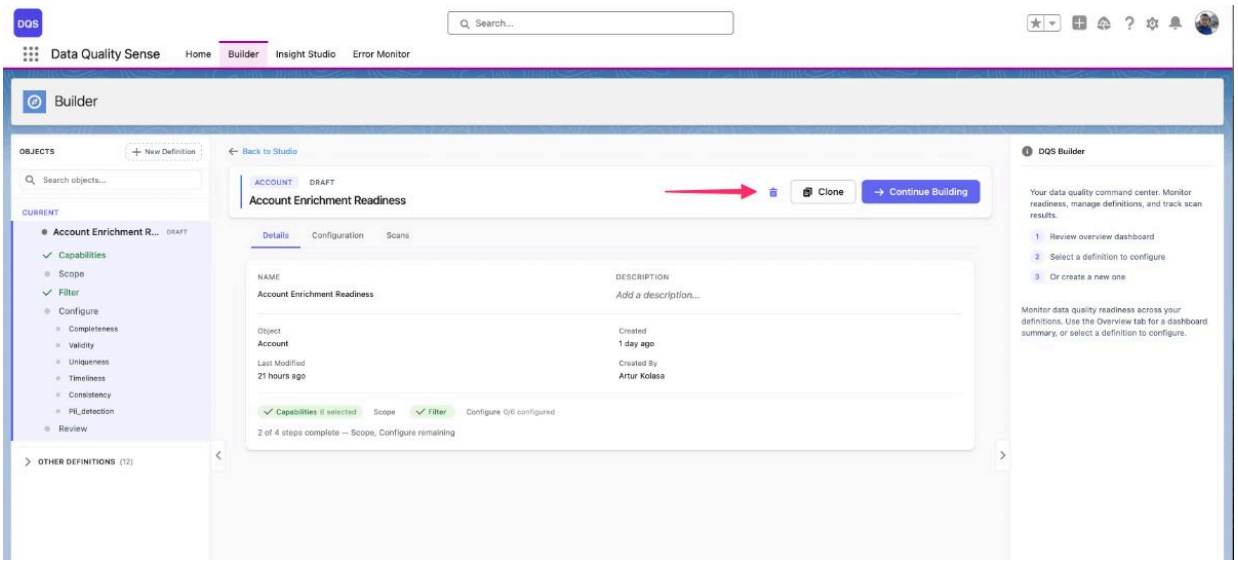

Activation Validation

[Section titled “Activation Validation”](#)

Before a definition can be activated, DQS validates:

Rule	Description
Fields selected	At least one field must be selected
Capabilities enabled	At least one capability must be enabled
Configuration complete	All required capability settings must be filled

ⓘNote

When you click **Activate**, the definition moves directly to Active status. The Ready state is used internally for validation gating.

Changing Status

[Section titled “Changing Status”](#)

- **Draft** → **Active**: Click “Mark as Complete” on the summary screen — choose between Complete Only, Complete & Schedule, or Complete & Run Scan
- **Active** → **Obsolete**: Use the definition settings to retire it
- **Obsolete** → **Active**: Reactivate from the definition settings

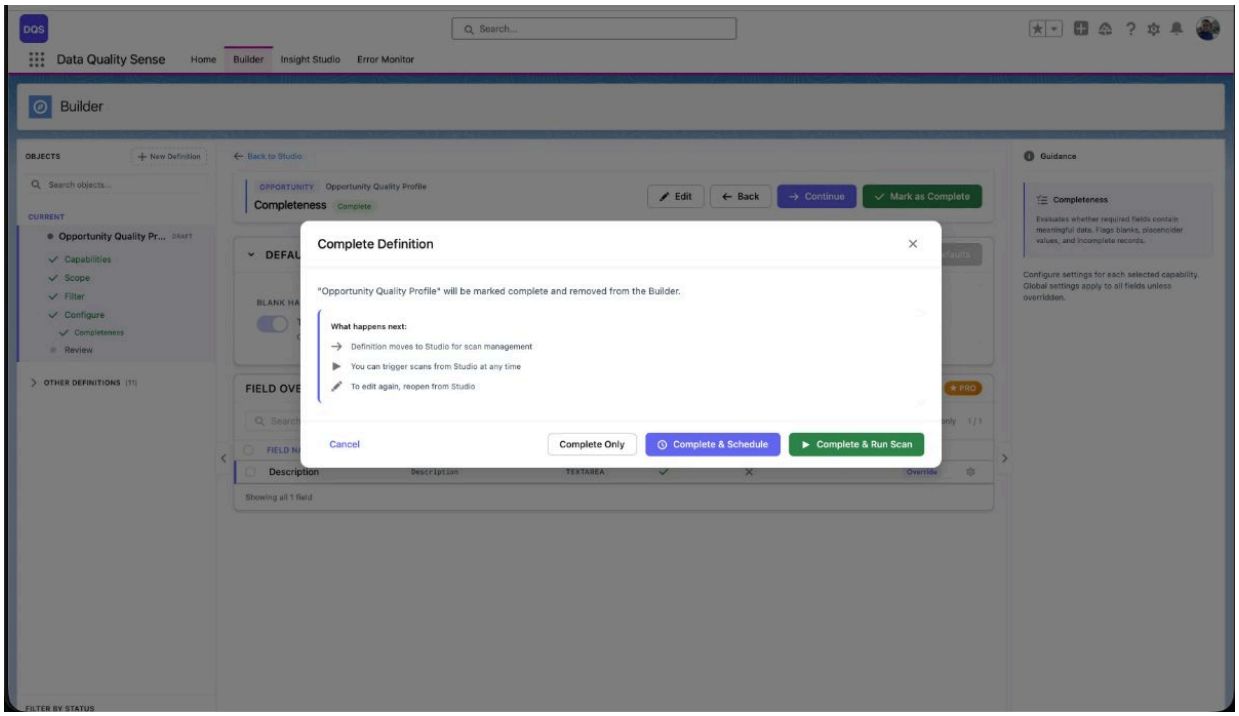
DQS
Search...
Data Quality Sense
Home
Builder
Insight Studio
Error Monitor
Builder
OBJECTS
New Definition
Search objects...
CURRENT
Opportunity Quality Pr... DRAFT
Capabilities
Scope
Filter
Configure
Completeness
Review
OTHER DEFINITIONS (11)
Back to Studio
OPPORTUNITY
Opportunity Quality Profile
Completeness
Complete
Edit
Back
Continue
Mark as Complete
Guidance
Completeness
Evaluates whether required fields contain meaningful data. Flags blanks, placeholder values, and incomplete records.
Configure settings for each selected capability. Global settings apply to all fields unless overridden.
Complete Definition
"Opportunity Quality Profile" will be marked complete and removed from the Builder.
What happens next:
Definition moves to Studio for scan management
You can trigger scans from Studio at any time
To edit again, reopen from Studio
Cancel
Complete Only
Complete & Schedule
Complete & Run Scan
Description
TEXTAREA
Showing all 1 field
FILTER BY STATUS

Capabilities

Quality Capabilities

The 6 Dimensions of Data Quality

Section titled “The 6 Dimensions of Data Quality”

Data Quality Sense evaluates your data across **6 distinct quality dimensions** (capabilities). Each dimension focuses on a different aspect of data quality and produces independent scores that roll up into an overall quality rating.

Completeness

Measures whether fields contain values. Detects null, blank, and missing data across your selected fields. Learn more →

Validity

Checks whether values conform to expected formats, ranges, and patterns. Supports picklist validation and regex matching. Learn more →

Uniqueness

Identifies duplicate values across records. Flags fields where unique values are expected but duplicates exist. Learn more →

Timeliness

Evaluates whether data is current and up-to-date. Measures freshness based on configurable time windows. Learn more →

Consistency

Checks logical consistency between related fields. Detects contradictions like a closed date before an open date. Learn more →

PII Detection

Scans for personally identifiable information in free-text fields. Helps with data privacy compliance. Learn more →

How Scoring Works

Section titled “How Scoring Works”

Each capability produces a **score from 0 to 100** for every scanned field:

- **100** — Perfect quality for this dimension
- **75–99** — Good, with minor issues
- **50–74** — Moderate quality, attention recommended
- **0–49** — Poor quality, action required
- **0** — No data to measure (e.g., all fields empty)

Scores are aggregated:

1. **Field Score** — Individual field result per capability
2. **Dimension Score** — Average across all fields for one capability
3. **Definition Score** — Weighted average across all dimensions

Capability Applicability

Section titled “Capability Applicability”

Not all capabilities apply to all field types. DQS automatically handles non-applicable combinations:

Field Type	Completeness	Validity	Uniqueness	Timeliness	Consistency	PII Detection
Text	✓	✓	✓	—	✓	✓
Number	✓	✓	✓	—	✓	—
Date	✓	✓	—	✓	✓	—
Picklist	✓	✓	—	—	✓	—
Boolean	✓	—	—	—	✓	—
Email	✓	✓	✓	—	—	✓
Phone	✓	✓	✓	—	—	✓

Completeness

What is Completeness?

Section titled “What is Completeness?”

Completeness measures the **fill rate** of your fields — the percentage of records where a field contains a non-null, non-blank value.

How It Works

Section titled “How It Works”

For each field included in the scan, the Completeness strategy:

1. Counts the total number of records in scope
2. Counts how many records have a non-empty value for that field
3. Calculates the fill rate: `(populated records / total records) × 100`

Configuration

Section titled “Configuration”

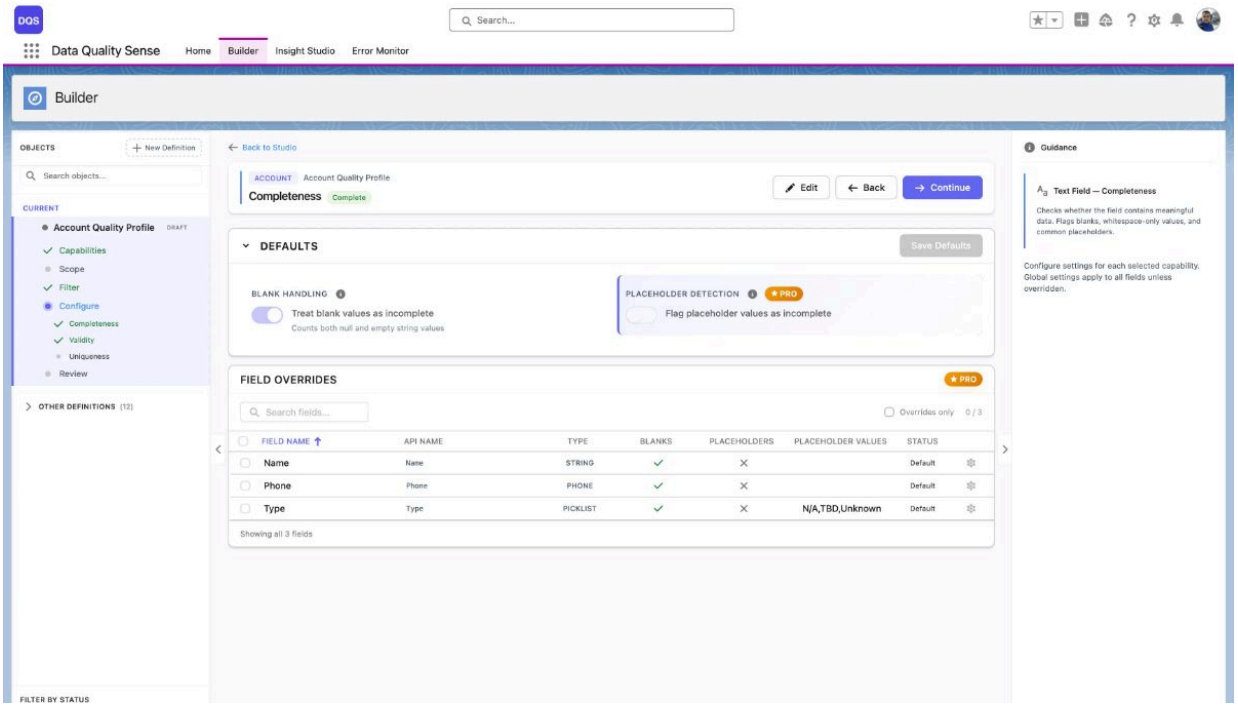

Global Settings

Section titled “Global Settings”

Setting	Description	Default
Expected Fill Rate	The minimum acceptable percentage of populated records	80%

Per-Field Overrides

Section titled “Per-Field Overrides”

Override the expected fill rate for individual fields. Common overrides:

- **Required fields** (Email, Name) → 100%
- **Optional fields** (Description, Notes) → 50%
- **Rarely used fields** → 20%

Scoring

Section titled “Scoring”

Fill Rate	Score
≥ Expected	100
Below expected	Proportional (e.g., 70% fill with 80% target = 87.5 score)
0%	0

ⓘNote

A score of 0 means the denominator was zero (no records to measure), not that all fields are empty.

Use Cases

Section titled “Use Cases”

- Identify fields that are frequently left blank by sales reps
- Monitor data entry compliance for required fields
- Track completeness trends over time after data quality initiatives
- Ensure critical fields (Email, Phone) are consistently populated

Bulk Configuration

Section titled “Bulk Configuration”

Use the **Bulk Config** option to set the same fill rate override across multiple fields at once — useful when you have many fields that share the same completeness requirement.

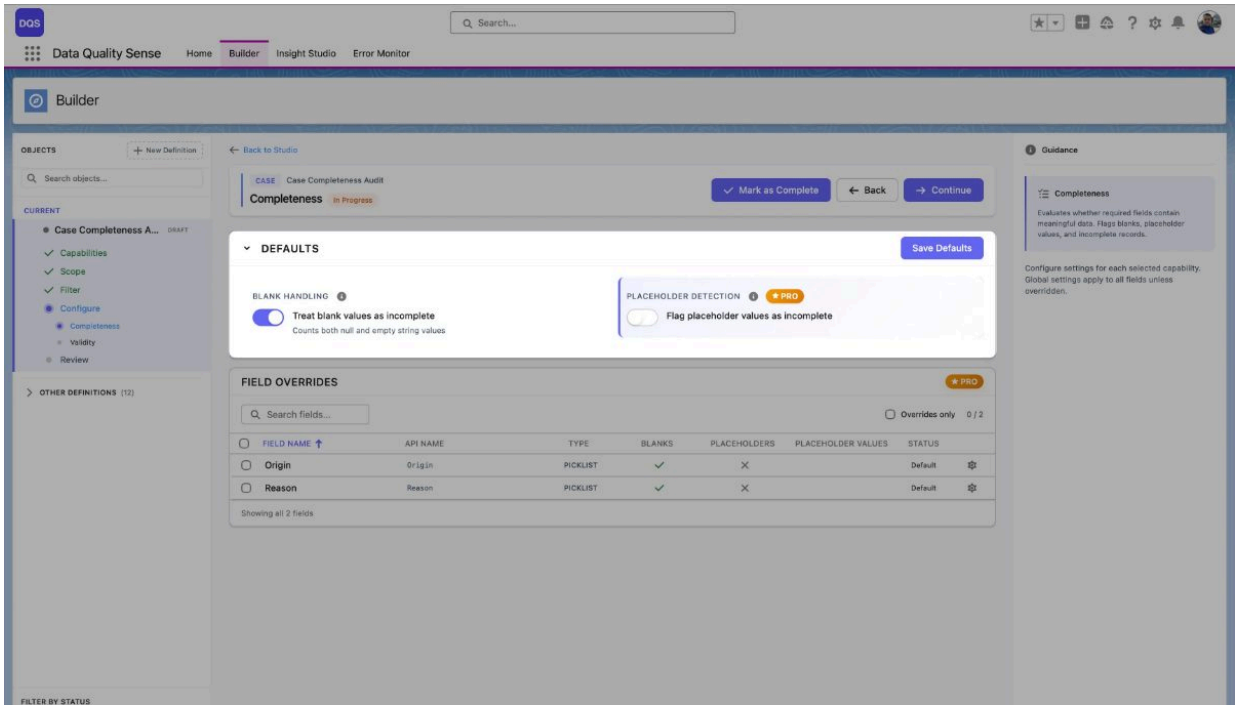

Validity

What is Validity?

Section titled “What is Validity?”

Validity measures whether field values **conform to expected formats, ranges, and patterns**. A field can be populated (complete) but still contain invalid data — Validity catches these issues.

How It Works

Section titled “How It Works”

The Validity strategy evaluates each field value against expected rules:

1. **Picklist fields** — Checks that values match the defined picklist values (including metadata and live values)
2. **Text fields** — Validates against format patterns (e.g., email format, phone format)
3. **Number fields** — Validates against expected ranges
4. **Date fields** — Checks for reasonable date ranges

Configuration

Section titled “Configuration”

Global Settings

Section titled “Global Settings”

The **Defaults** section controls global validation options that apply to all fields:

Setting	Description
Include blank values in validation	When enabled, blank/null values will fail validation
Case-sensitive matching	Pattern matching considers uppercase/lowercase
Pattern Selection	Choose a default validation pattern (Email, URL, Fixed Length, or Custom regex)

The **Field Overrides** table below lists each field in scope with its current pattern and status. Fields marked “Default” use the global settings, “None” means no pattern is assigned yet.

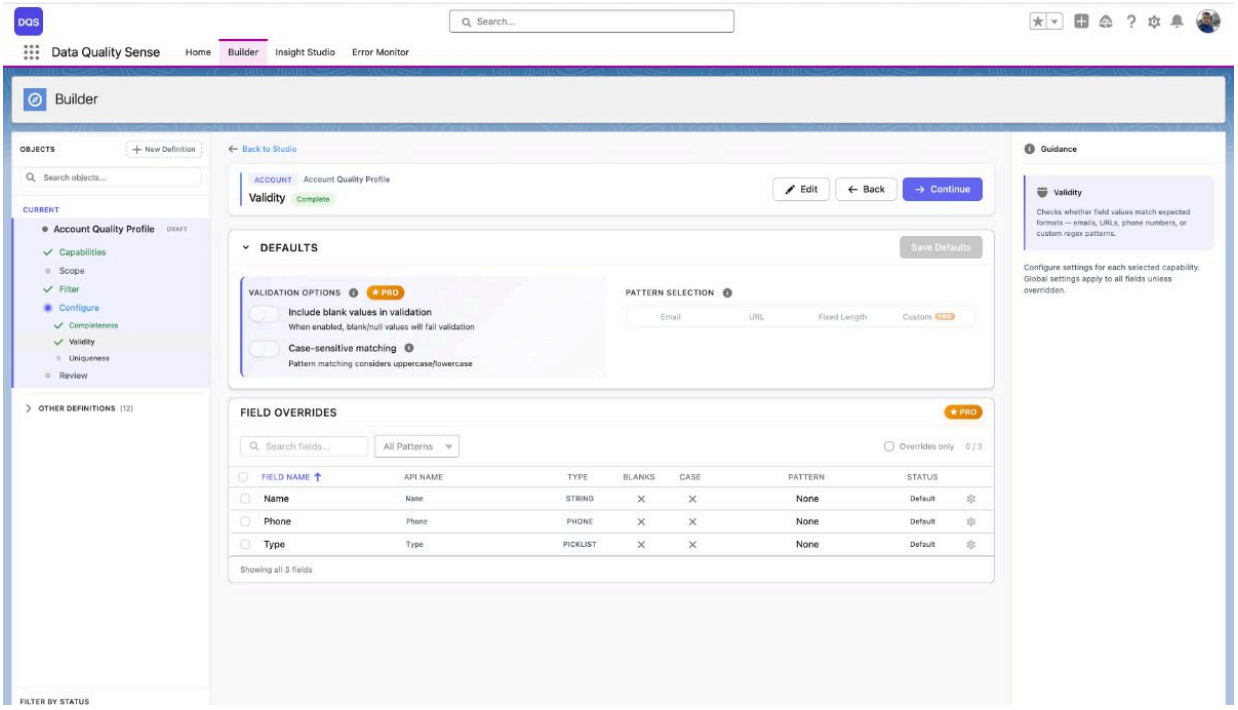

Per-Field Overrides

[Section titled “Per-Field Overrides”](#)

Click on a field in the Field Overrides table to open its configuration modal. Here you can assign a specific validation pattern for that field — choose from predefined patterns (Email, URL, Fixed Length) or select **Custom** to enter your own regex. Each field override also lets you toggle **Include blank values** and **Case-sensitive matching** independently from the global defaults. Use the **Revert to Global** link to reset the field back to the global settings.

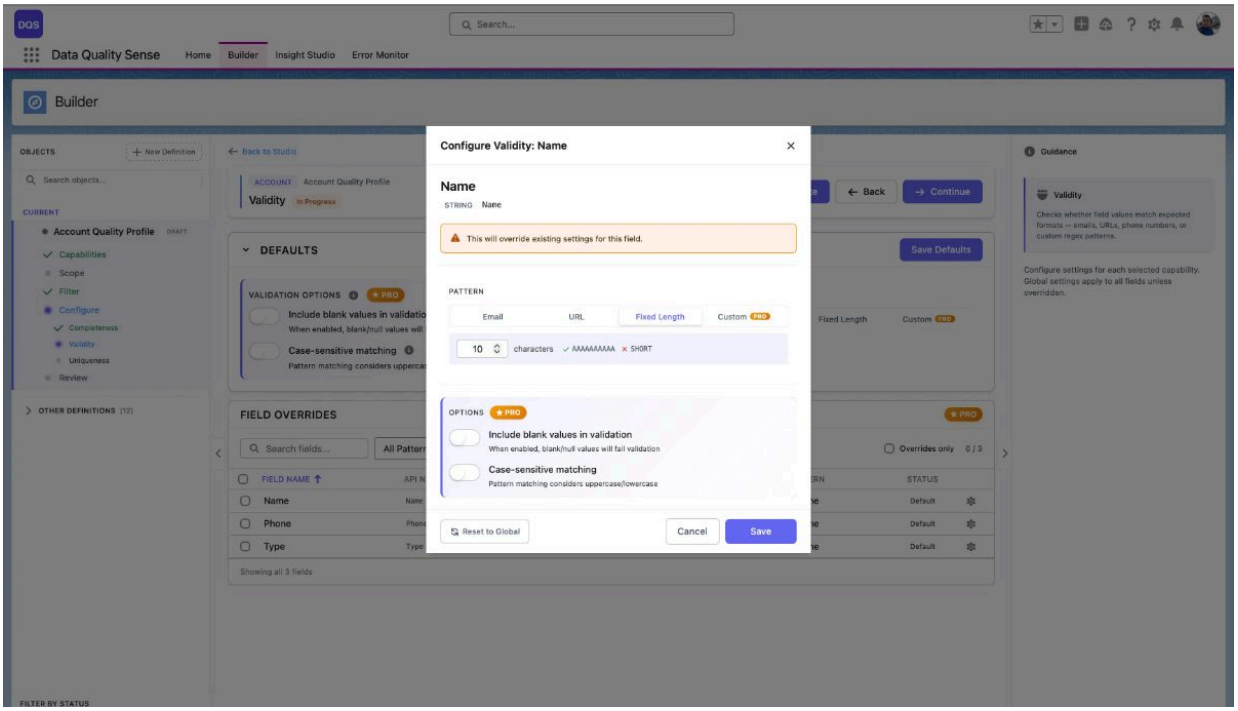

Scoring

Section titled “Scoring”

Result	Score
All values valid	100
Some invalid	Proportional to valid percentage
All invalid	0
No data	0

Regex Patterns

Section titled “Regex Patterns”

DQS uses **Java-compatible regular expressions** for text field validation. When you select **Custom** in the pattern picker, a text field appears where you can enter your own regex pattern.

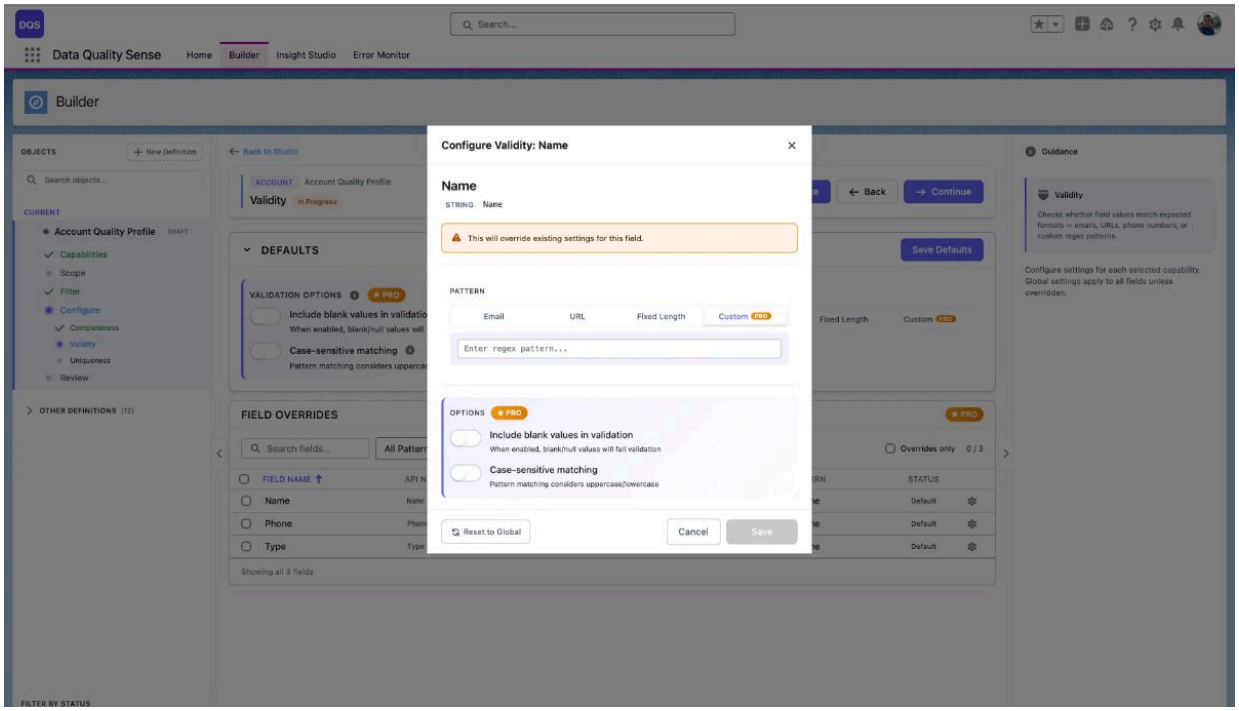

See the Regex Tester for an interactive tester and a library of ready-to-use patterns for email, phone, URL, postal codes, and more.

Use Cases

Section titled “Use Cases”

- Ensure email fields contain properly formatted email addresses
- Verify picklist fields only contain approved values
- Detect free-text entries in fields that should use controlled vocabularies
- Validate phone number formats across regions

Uniqueness

What is Uniqueness?

Section titled “What is Uniqueness?”

Uniqueness measures whether field values are **distinct across records**. High uniqueness means each record has a different value for the field — low uniqueness indicates duplicates.

How It Works

Section titled “How It Works”

For each field, the Uniqueness strategy:

1. Collects all non-null values across records in scope
2. Identifies duplicate values
3. Calculates: `(unique values / total populated values) × 100`

Configuration

Section titled “Configuration”

Global Settings

Section titled “Global Settings”

The **Defaults** section controls global uniqueness options:

Setting	Description
Case-sensitive matching	When enabled, “John Smith” and “john smith” are considered different values for comparison. When disabled, they count as duplicates.
Include blanks in uniqueness checks	When enabled, blank and null values are treated as distinct values in comparison calculations.

The **Field Overrides** table below lists each field with its current Case Sensitive, Include Blanks settings, and status.

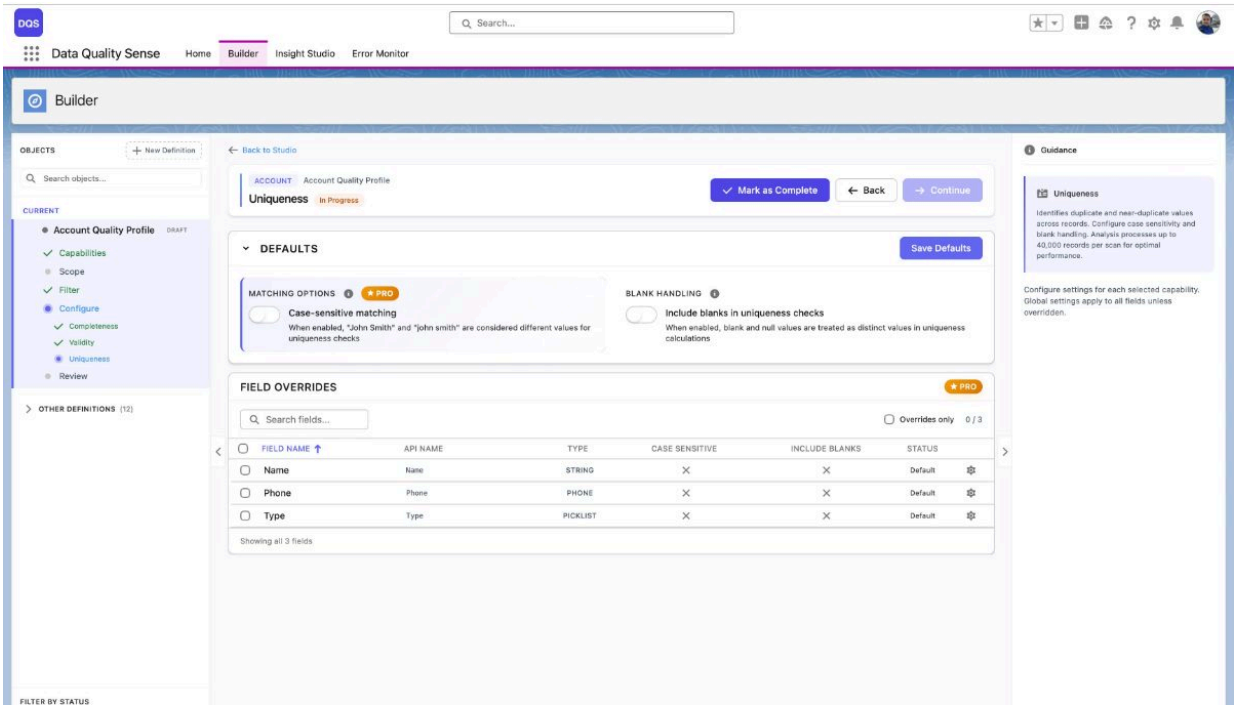

Per-Field Overrides

[Section titled “Per-Field Overrides”]

Click on a field in the Field Overrides table to open its configuration modal. You can toggle **Case-sensitive matching** and **Include blanks in uniqueness checks** independently from the global defaults. Use the **Revert to Global** link to reset the field back to global settings.

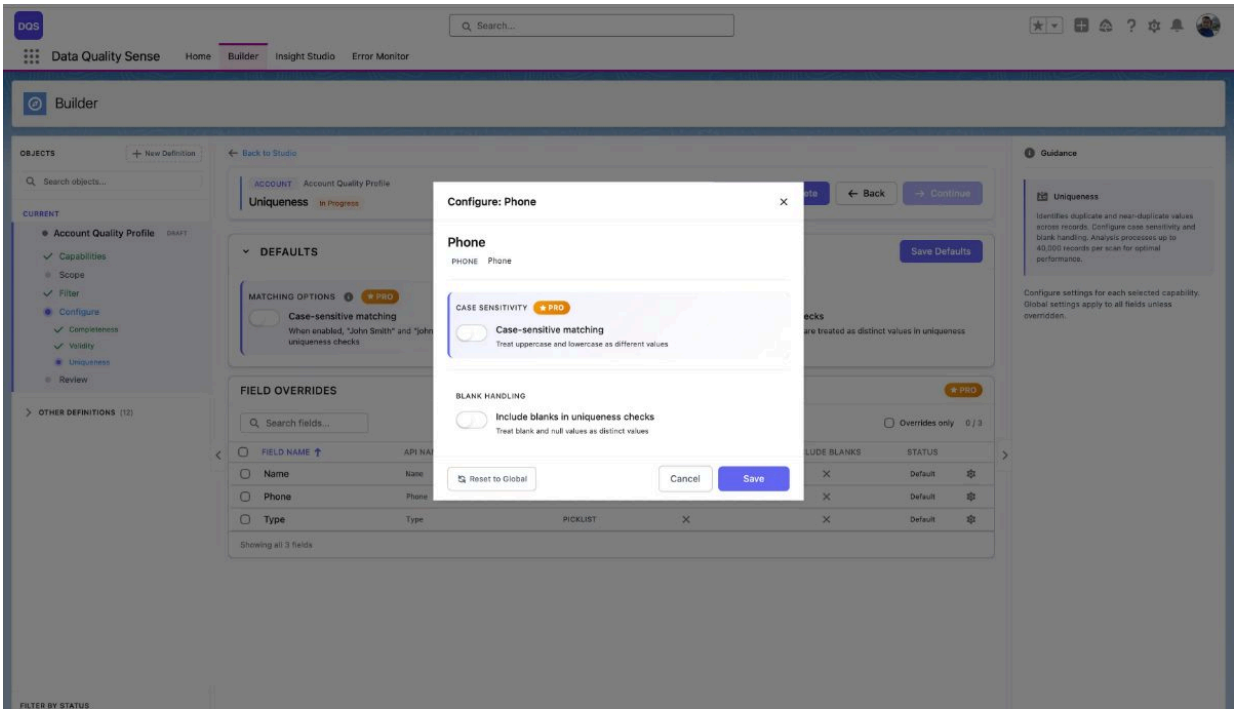

Scoring

[Section titled “Scoring”]

Result	Score
All values unique	100
Some duplicates	Proportional to unique percentage
All values identical	Near 0
No data	0

Analysis Limit

Section titled “Analysis Limit”

Uniqueness analysis processes up to **40,000 records per scan**. For objects with more records, results reflect a representative sample. This limit exists to prevent Salesforce heap memory overflow, since the engine builds an in-memory map of value counts per field. Fields that exceed 40,000 distinct values are flagged as high cardinality fields.

Applicable Field Types

Section titled “Applicable Field Types”

Uniqueness is most meaningful for:

- **Email** — should be unique per contact/lead
- **Phone** — often unique per person
- **External IDs** — must be unique by definition
- **Text fields** — names, descriptions

Less meaningful for:

- **Boolean** — only two possible values
- **Picklist** — limited value set by design
- **Date** — many records may share dates

Use Cases

Section titled “Use Cases”

- Detect duplicate email addresses across Contacts or Leads
- Verify external ID fields are truly unique
- Identify data entry issues where the same value is copied across records
- Monitor deduplication efforts over time

Timeliness

What is Timeliness?

Section titled “What is Timeliness?”

Timeliness measures whether date and datetime fields contain **recent, up-to-date values**. It answers the question: “Is this data still fresh?”

How It Works

Section titled “How It Works”

For each date field, the Timeliness strategy:

1. Reads the date/datetime value
2. Calculates the age (difference between the value and the current date)
3. Compares the age against the configured freshness window
4. Marks values older than the window as “stale”

Configuration

Section titled “Configuration”

Global Settings

Section titled “Global Settings”

The **Defaults** section controls global timeliness options:

Setting	Description
Freshness Threshold	The Default Freshness Window in days — records older than this threshold are considered stale.
Null Handling	**Treat null dates as stale** — when enabled, records with no date value count as stale data.
Overdue Tracking	Enable overdue tracking by default. Flags records past their expected date. *(PRO)*
Operational Range	Enable operational range validation by default. Checks whether dates fall within an acceptable time span. *(PRO)*

The **Field Overrides** table lists each date field with its current threshold, overdue, and window settings.

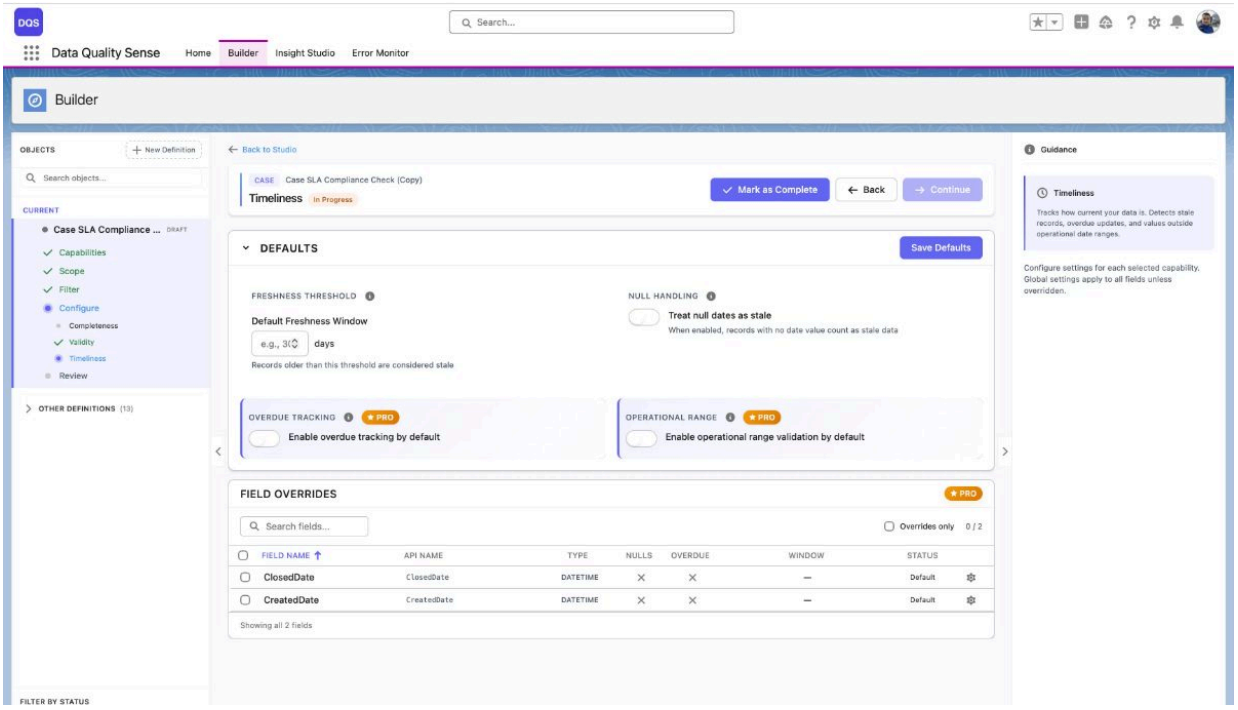

Per-Field Overrides

[Section titled “Per-Field Overrides”]

Click on a field in the Field Overrides table to open its configuration modal. You can set a custom **Freshness Threshold** (in days), toggle **Null Handling**, **Overdue Tracking**, and **Operational Range** independently from the global defaults. Use the **Revert to Global** link to reset the field back to global settings.

Different date fields may have different freshness requirements:

Field Example	Recommended Window
`LastActivityDate`	7 days
`LastModifiedDate`	30 days
`Contract_End_Date__c`	90 days
`Annual_Review_Date__c`	365 days

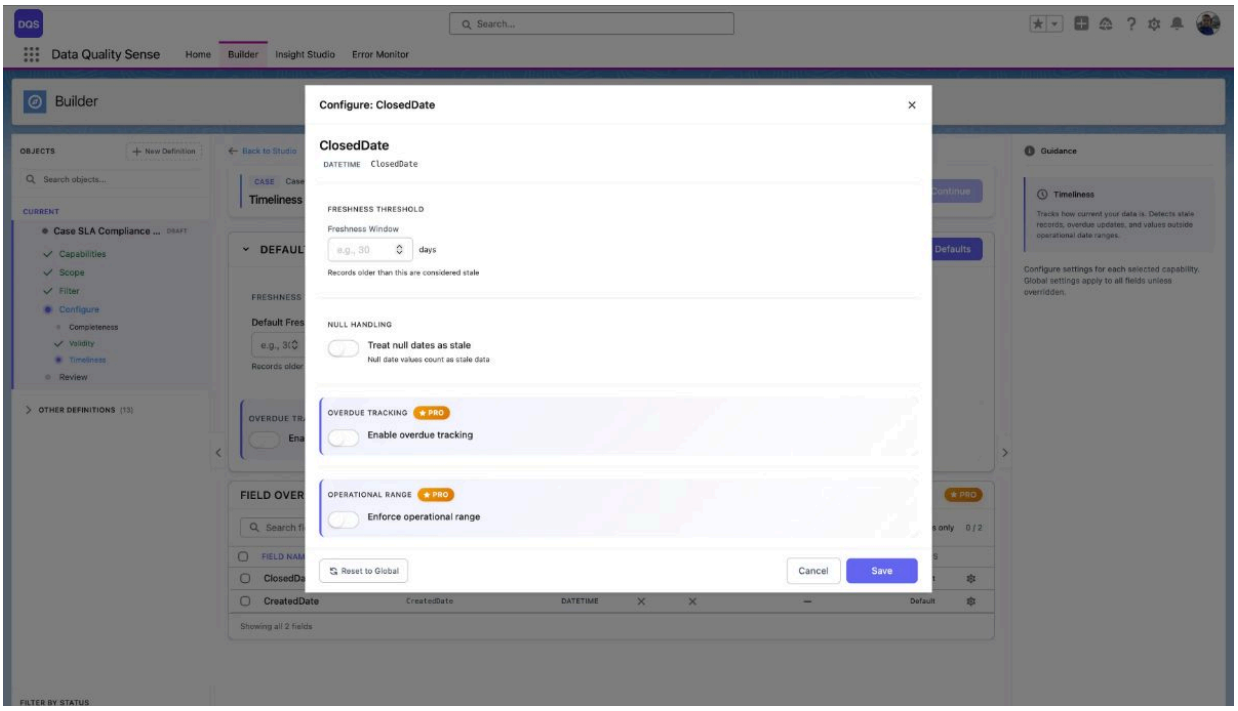

Scoring

[Section titled “Scoring”](#)

Result	Score
All dates within window	100
Some stale dates	Proportional to fresh percentage
All dates stale	0
No data	0

Use Cases

[Section titled “Use Cases”](#)

- Monitor whether sales reps are keeping opportunity dates current
- Track stale contacts that haven’t been updated recently
- Ensure compliance-related dates are being maintained
- Identify “dead” records that may need cleanup or archival

Consistency

What is Consistency?

Section titled “What is Consistency?”

Consistency measures whether **related fields contain logically compatible values**. Data can be complete and valid individually, but still be inconsistent when fields contradict each other.

How It Works

Section titled “How It Works”

The Consistency strategy evaluates relationships between pairs or groups of fields:

1. Identifies configured consistency rules (field relationships)
2. For each record, checks whether the related fields satisfy the rule
3. Calculates the percentage of records passing all consistency checks

Configuration

Section titled “Configuration”

Global Settings

Section titled “Global Settings”

The **Defaults** section controls global consistency options:

Setting	Description
Default Allowed Values	No defaults are set globally — configure allowed values per field to define what’s considered consistent.
Case-sensitive matching	When enabled, “Active” and “active” are treated as different values. Disabled by default.
Top N Values	*(PRO)* Analyze only the top N most frequent values for consistency checks.
Minimum Frequency	*(PRO)* Ignore values that appear fewer times than this threshold.

The **Field Overrides** table lists each field with its current allowed values source, case sensitivity setting, and status.

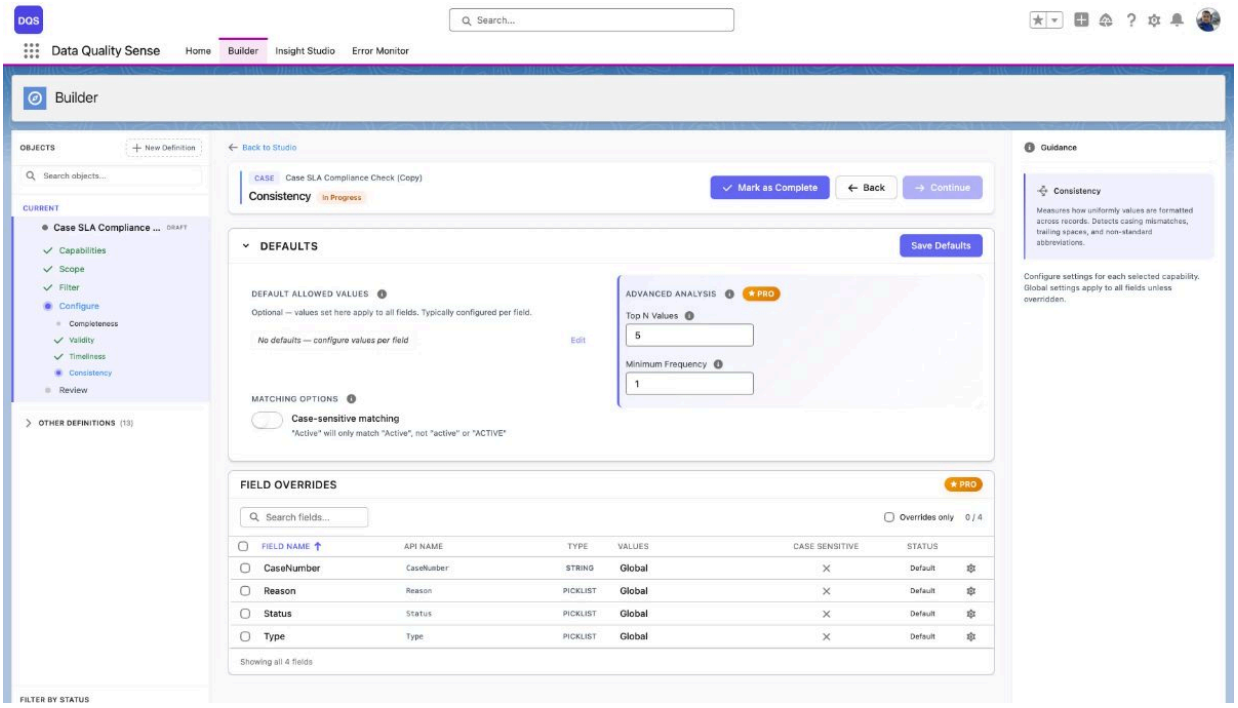

Per-Field Overrides

[Section titled “Per-Field Overrides”](#)

Click on a field in the Field Overrides table to open its configuration modal. You can define **allowed values** for that field by adding them manually or importing from field metadata. Toggle **Case-sensitive matching** and configure **Advanced Analysis** settings (Top N Values, Min Frequency) independently from the global defaults. Use the **Revert to Global** link to reset.

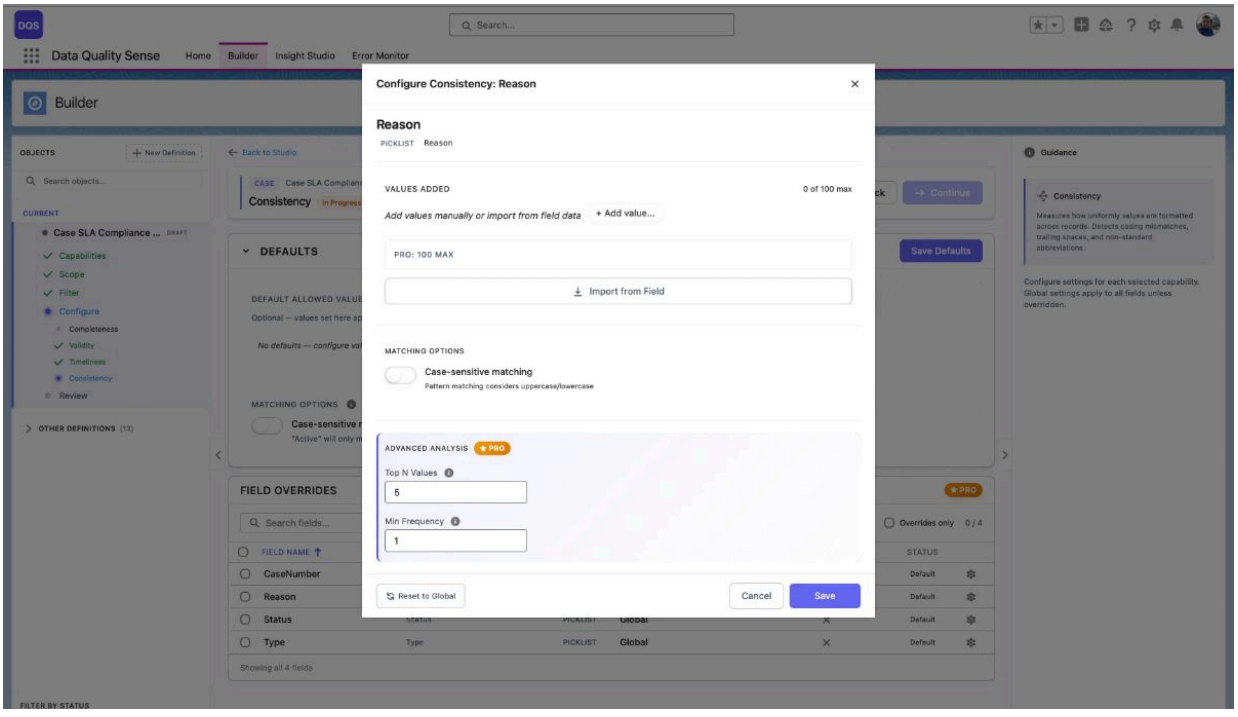

Import from Field

[Section titled “Import from Field”](#)

For picklist fields, click **Import from Field** to load all existing values directly from the field metadata. The import dialog shows each value with a checkbox and the number of records using it, so you can select which values to treat as valid. Click **Apply Selection** to confirm.

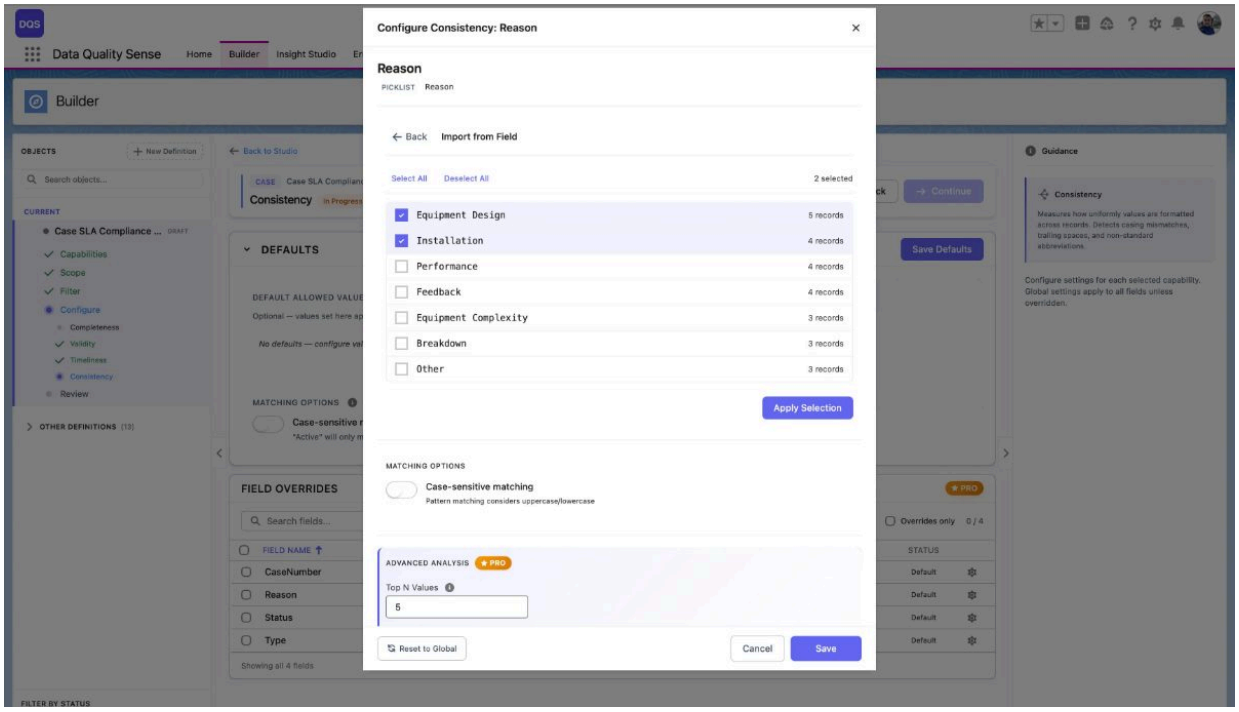

Scoring

Section titled “Scoring”

Result	Score
All records consistent	100
Some inconsistencies	Proportional to consistent percentage
All records inconsistent	0
No data	0

Use Cases

Section titled “Use Cases”

- Detect opportunities where Close Date is before Created Date
- Find accounts where Billing and Shipping addresses contradict
- Identify leads where Country doesn’t match the phone number format
- Monitor cross-field data entry quality

PII Detection

What is PII Detection?

Section titled “What is PII Detection?”

PII Detection scans free-text fields for **personally identifiable information** that shouldn’t be stored in those fields. It helps organizations comply with data privacy regulations like GDPR, CCPA, and HIPAA.

How It Works

Section titled “How It Works”

The PII Detection strategy analyzes text content for patterns that match:

1. **Social Security Numbers** (SSN) — numeric patterns like XXX-XX-XXXX
2. **Credit Card Numbers** — 13–19 digit sequences matching card network patterns
3. **Email Addresses** — in fields where email storage is not expected
4. **Phone Numbers** — in free-text fields (not dedicated phone fields)
5. **National ID Numbers** — country-specific patterns

Configuration

Section titled “Configuration”

Detection Patterns

Section titled “Detection Patterns”

The **Defaults** section provides three preset groups of detection patterns:

Preset	Description
Standard	Core PII patterns — Social Security Number, Credit Card Number, Email Address, US Phone Number
Critical	High-risk financial and identity patterns
Extended	Full set including IP Address, IBAN, Date of Birth, International Phone, and more

Each pattern shows its regex expression and can be individually enabled or disabled. You can also add your own patterns in the **Add Custom Pattern** section by entering a regex and a label.

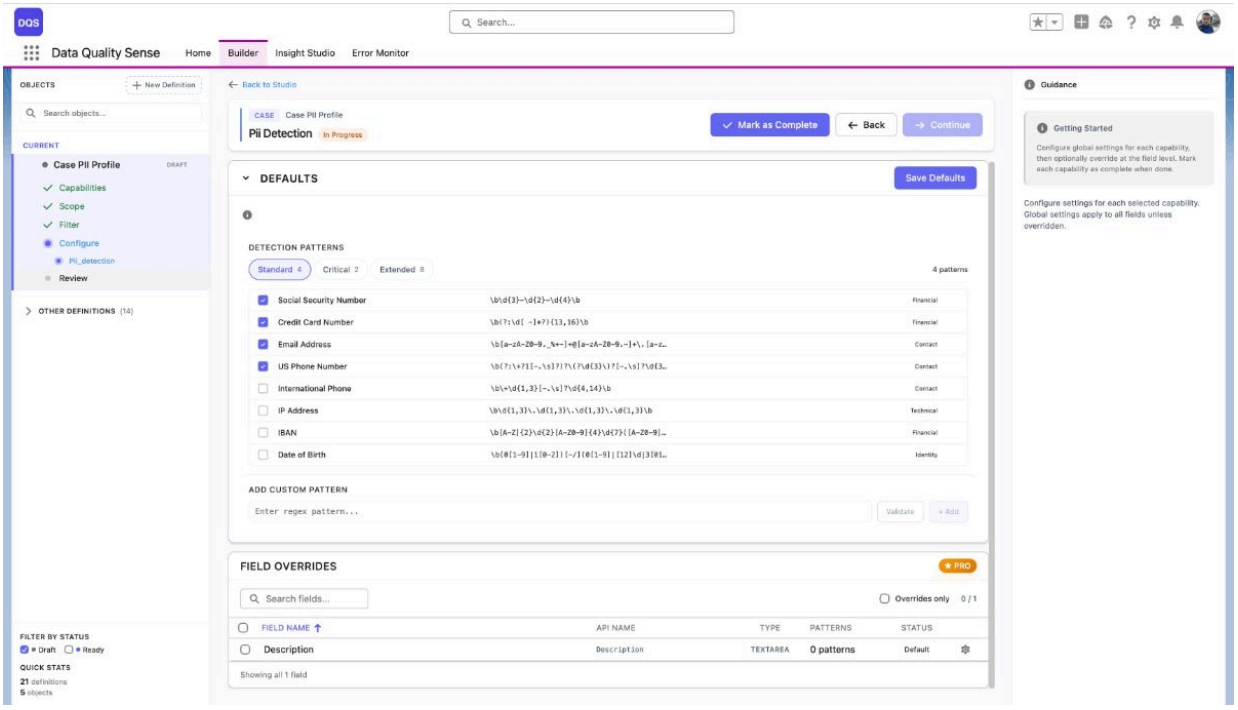

Per-Field Overrides

[Section titled “Per-Field Overrides”]

Click on a field in the Field Overrides table to open its configuration modal. You can select which detection patterns apply to that field — choose a preset or enable/disable individual patterns. The modal also lets you add **custom patterns** specific to that field. Use the **Revert to Global** link to reset.

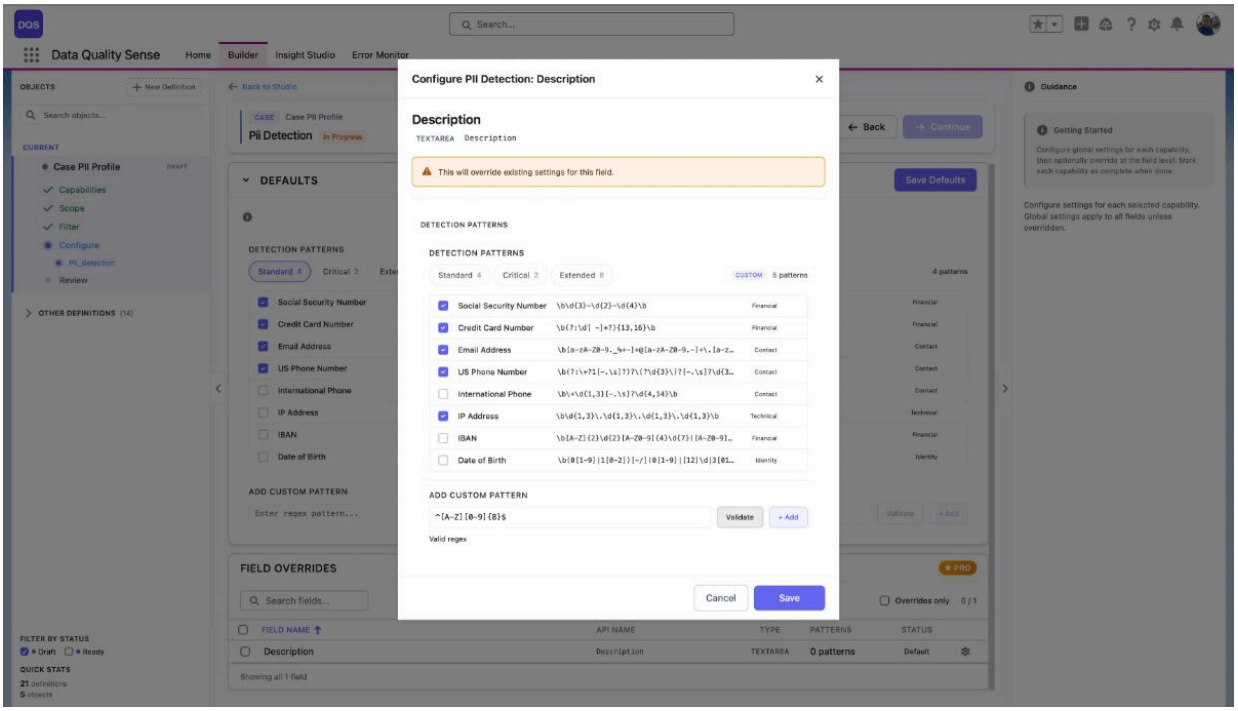

Scoring

[Section titled “Scoring”]

Result	Score
No PII detected	100
Some PII found	Proportional to clean percentage
PII in all records	0
No data	0

⚠Caution

PII Detection uses pattern matching, not AI classification. It may produce false positives (e.g., a number that looks like an SSN but isn't). Review flagged records before taking action.

PII Regex Patterns

Section titled "PII Regex Patterns"

DQS uses **Java-compatible regular expressions** to detect PII in free-text fields. See the Regex Tester for an interactive tester and a full library of PII patterns — including SSN, credit cards (Visa, Mastercard, Amex), IBAN, passport numbers, PESEL, NIP, and more.

Use Cases

Section titled "Use Cases"

- Audit Description and Notes fields for accidentally stored credit card numbers
- Monitor free-text fields for social security numbers
- Compliance reporting for data privacy regulations
- Pre-migration data cleansing

Insight Studio Overview

What is Insight Studio?

Section titled “What is Insight Studio?”

Insight Studio (DIS) is the **visualization and analytics layer** of Data Quality Sense. It consumes scan results and presents them as interactive dashboards, charts, and recommendations.

Workspace Layout

Section titled “Workspace Layout”

Insight Studio uses a **3-zone layout**:

- **Sidebar** (orange) — Object/definition navigation, filters, and breakdown options
- **Stage** (green) — Main content area with scores, charts, and data tables
- **Portfolio Overview** (red) — Summary metrics, quick links, and contextual actions

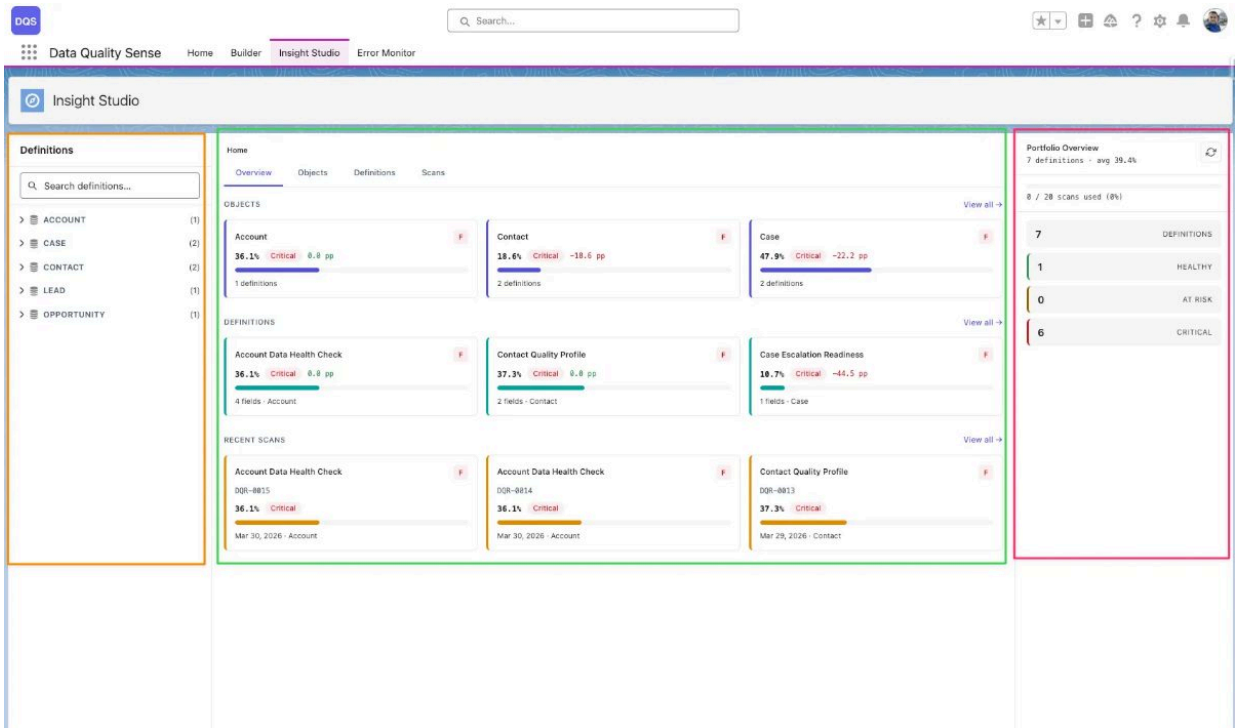

Score Overview

At-a-glance quality scores for each dimension, with overall ratings and grade indicators.

Trend Analysis

Sparklines and trend charts showing how data quality changes over time across scans.

Field Health

Matrix view of every field’s quality score per dimension. Quickly spot the weakest fields.

AI Mentor

Contextual recommendations based on scan results. Suggests actions to improve data quality.

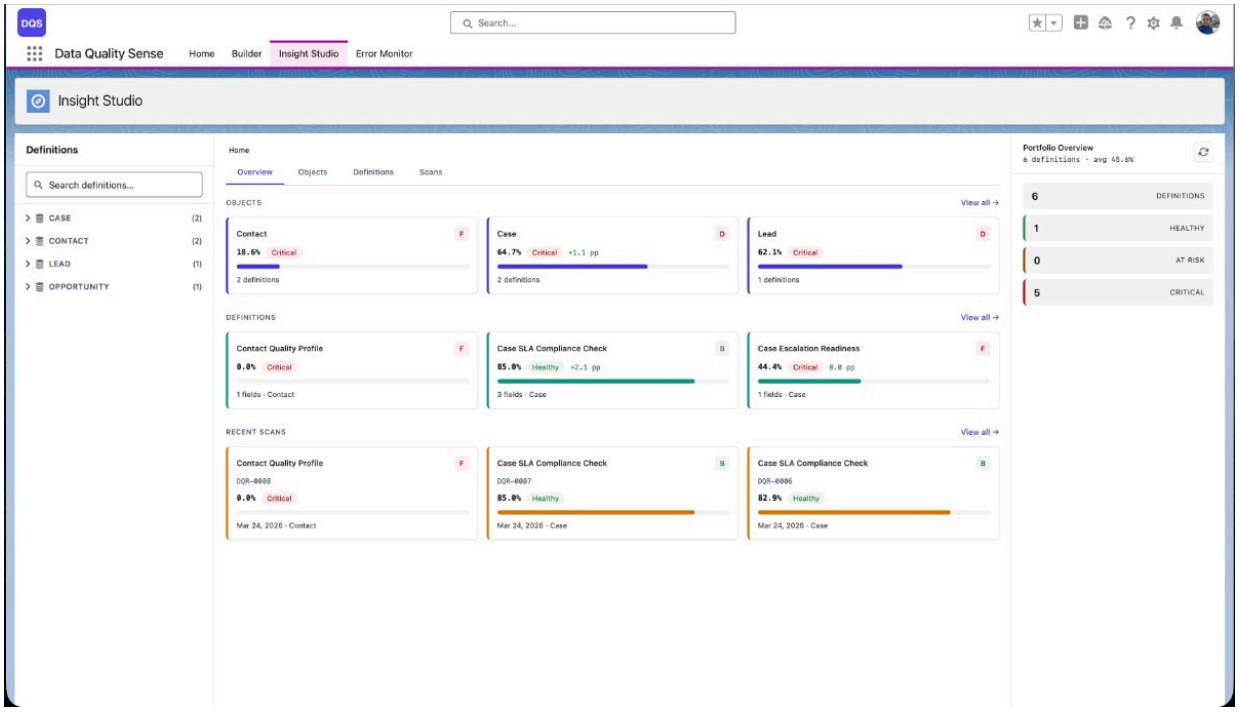

Key Features

Section titled “Key Features”

Feature	Description
Multi-level navigation	Drill from Home → Object → Definition → Scan → Field → Dimension
Score comparison	Compare results between two scans to see improvement or regression
Actions menu	Create Tasks, Post Chatter, or Export CSV for impacted records
Scan trigger	Manually trigger a scan from the dashboard
Schedule management	Create and manage scan schedules

Learn More

Section titled “Learn More”

- Navigation
- Scores & Trends
- Field Health
- Actions
- Exports

Navigation

Navigation Depth

Section titled “Navigation Depth”

Insight Studio provides a **6-level drill-down** hierarchy:

```
Home → Object → Definition → Scan → Field → Dimension
```

Level 1: Home

Section titled “Level 1: Home”

The top-level view showing all scanned objects with their latest quality scores. Use this to identify which objects need attention.

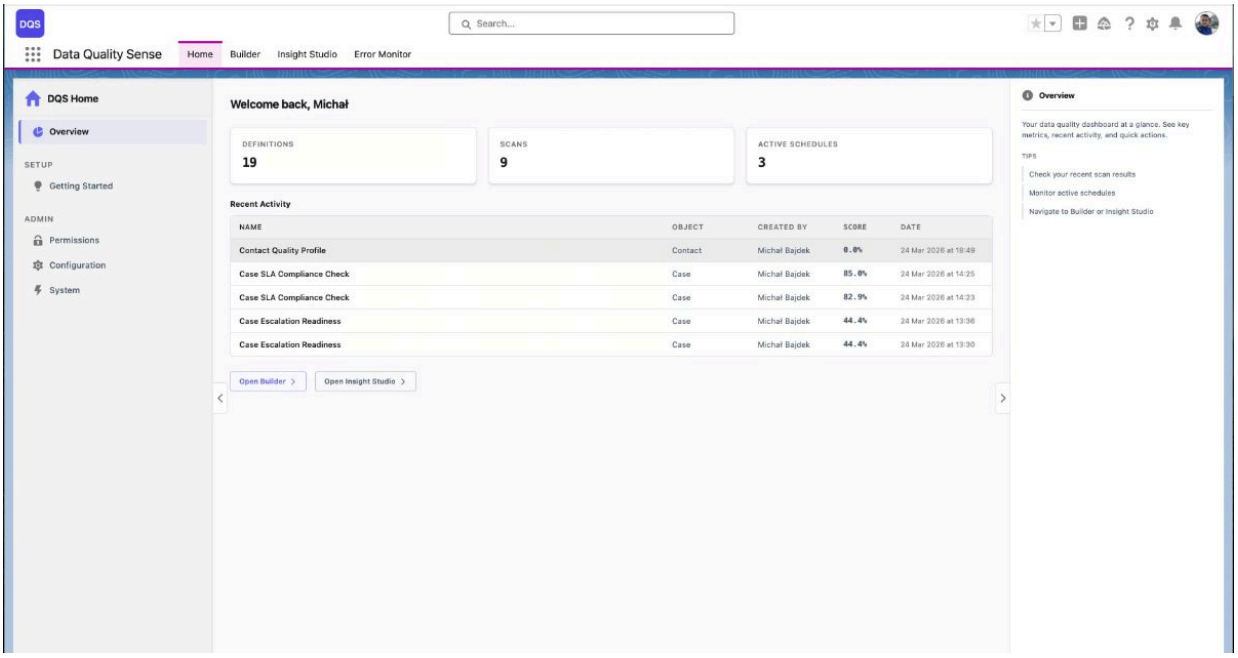

The Home page also includes a **Configuration** panel where administrators can edit default parameters such as Scan Result Retention Days, Purge Batch Size, Purge CRON Expression, and Error Log Retention Days directly from the UI.

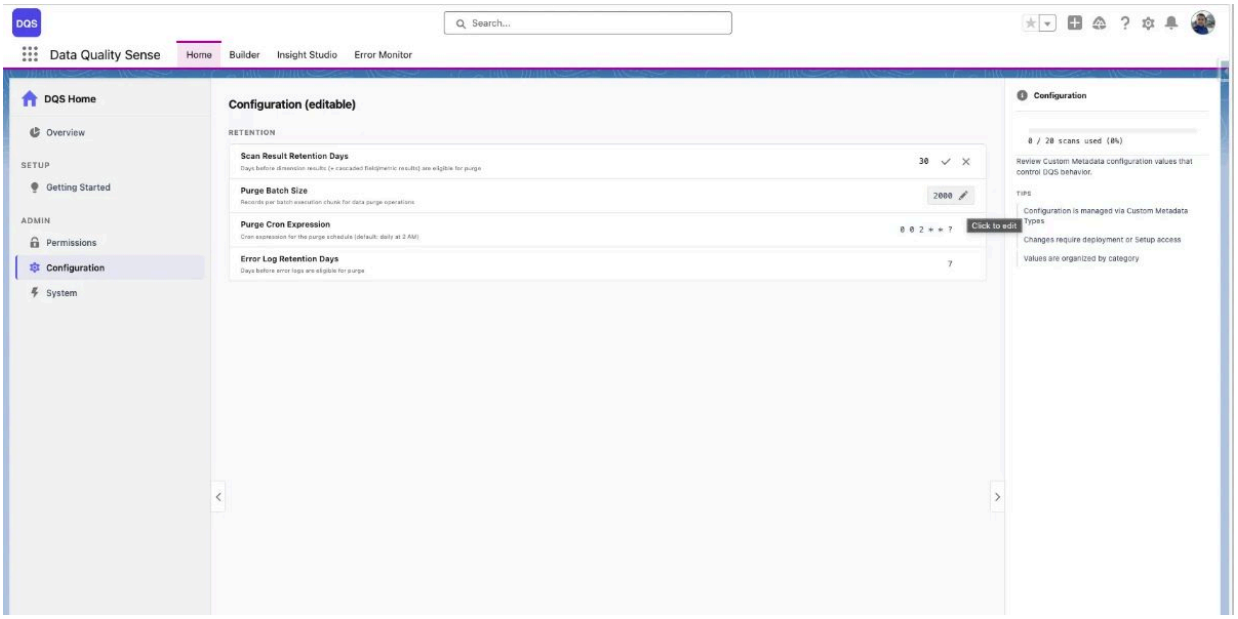

Level 2: Object

[Section titled “Level 2: Object”]

Shows all definitions for a selected object with the overall quality score, last scan date, status, and definition cards. Compare different scan configurations and their results side by side.

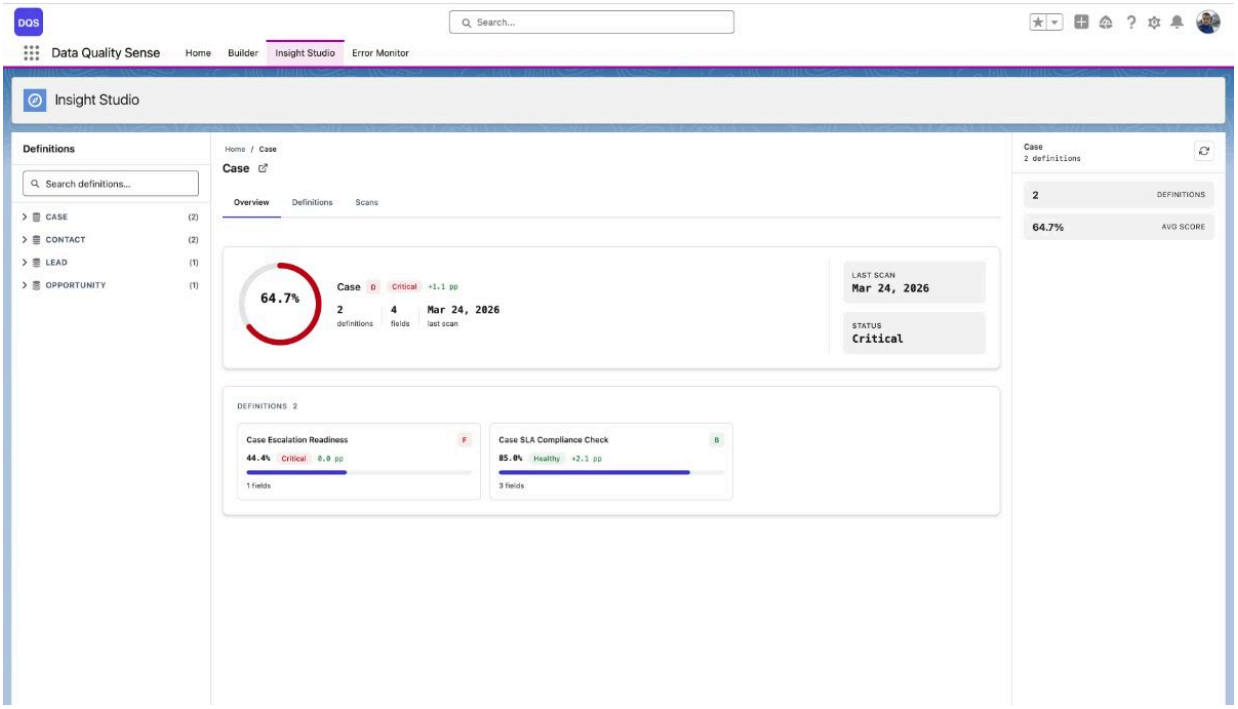

Level 3: Definition

[Section titled “Level 3: Definition”](#)

The main dashboard for a single definition. It has four tabs:

- **Overview** — Overall quality score, dimension score bars, Score Trend chart, and Record Scores panel
- **Fields** — Field Health Matrix with per-field scores across dimensions
- **Scans** — Full scan history with per-run results
- **Exports** — Download violation details as CSV per dimension and field, or take action via Create Tasks and Post Chatter

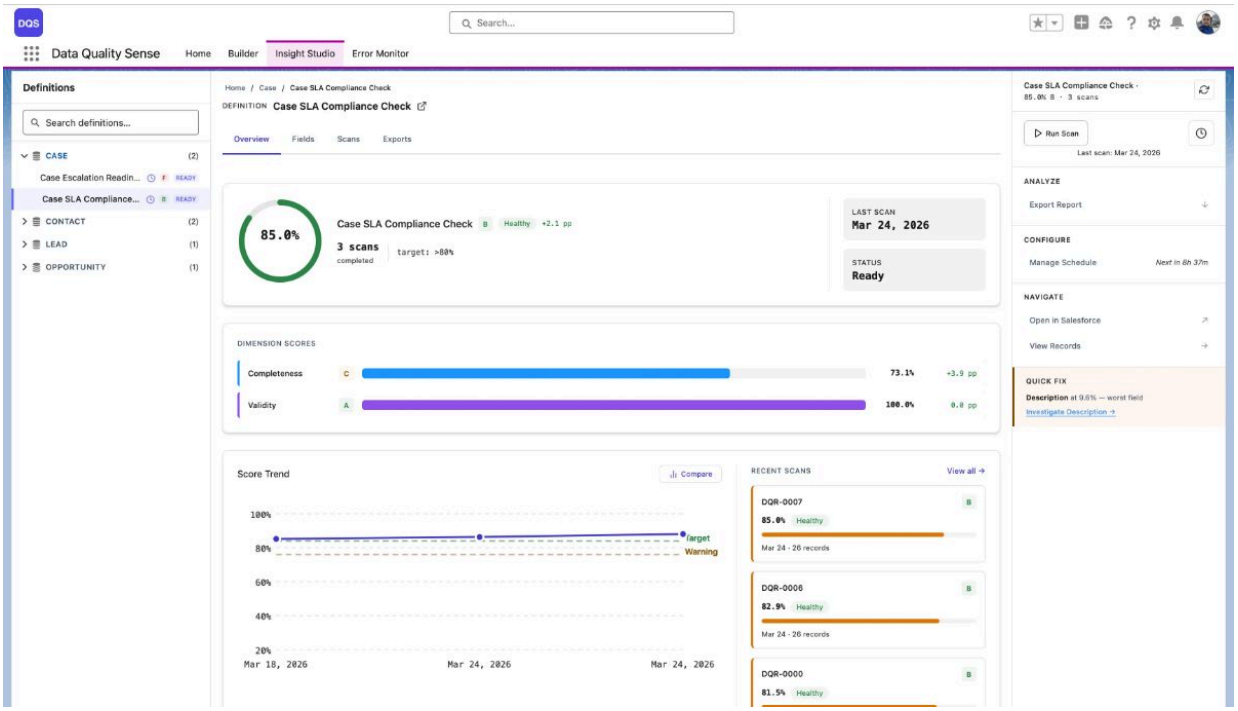

The **Fields** tab shows each field's score per dimension in a color-coded matrix — red for poor, green for good. Use it to quickly spot the weakest fields.

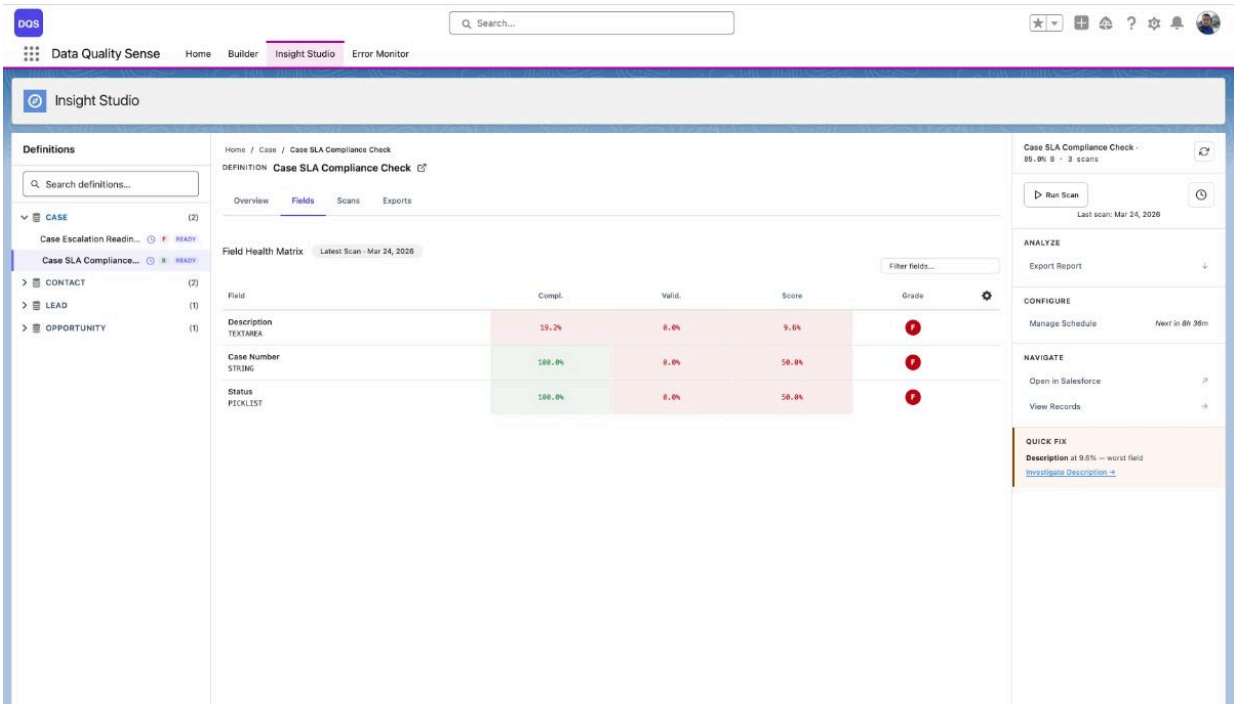

The **Exports** tab lists available exports per dimension and field, with download buttons and export history.

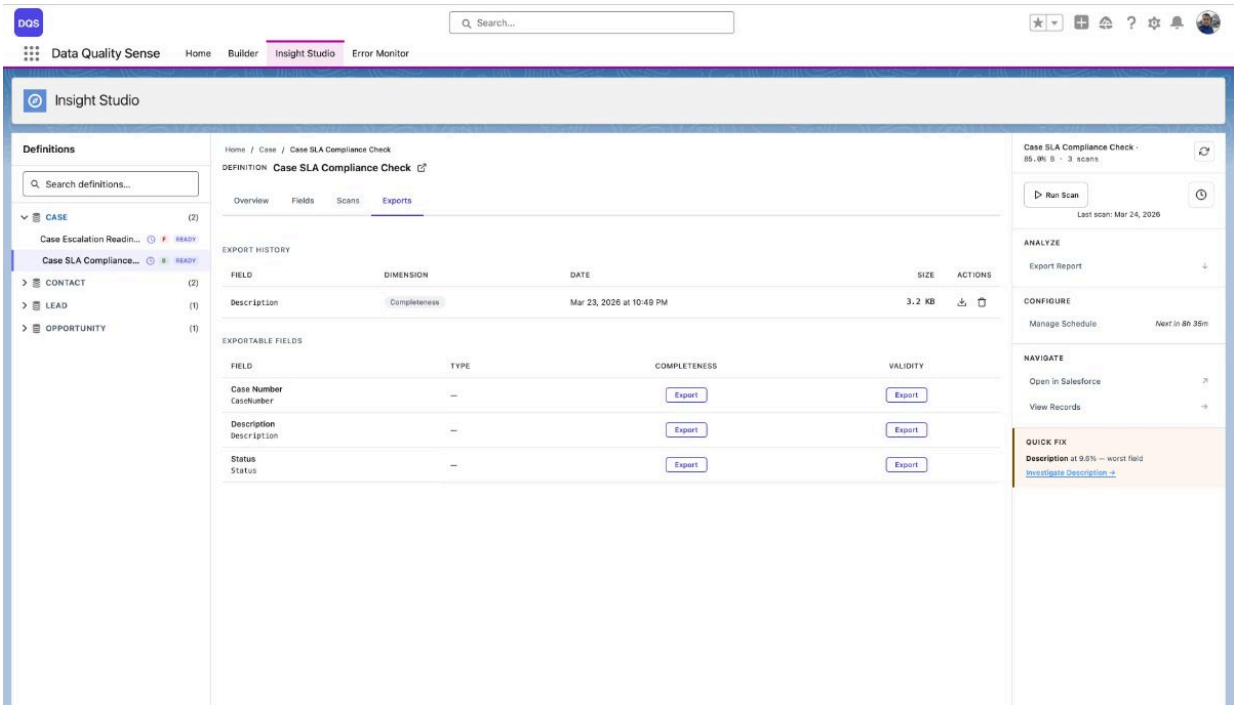

From the Overview tab, click **Compare** on the Score Trend chart to compare two scans side by side — see dimension-level and field-level deltas between a baseline and comparison scan.

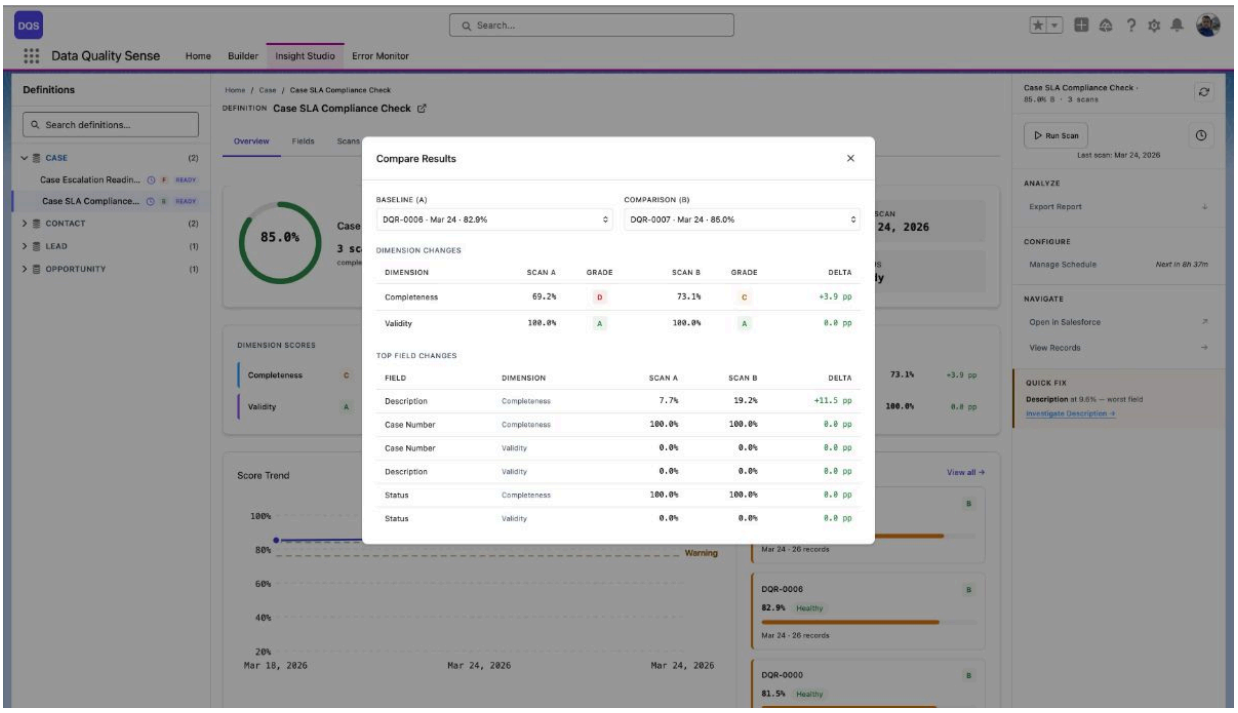

Level 4: Scan

[Section titled “Level 4: Scan”]

Detailed results for a specific scan execution. The **Scans** tab shows the full scan history with status indicators, scores, and per-dimension breakdowns for each run.

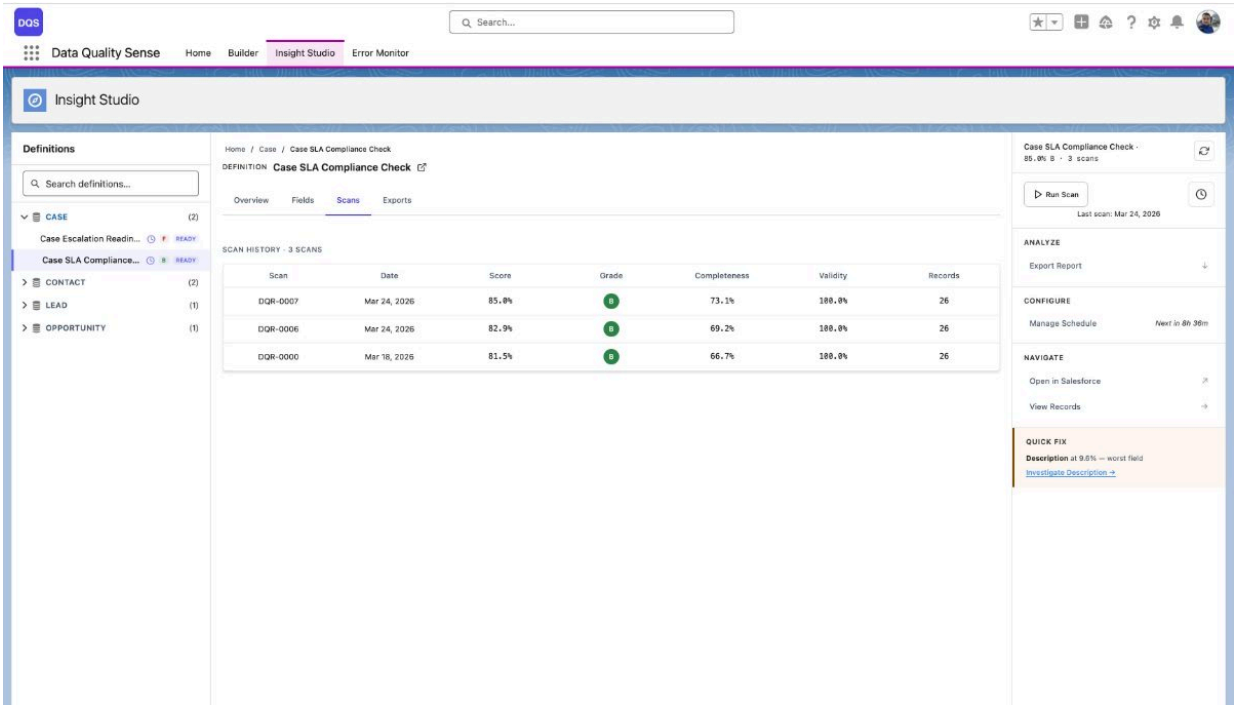

Level 5: Field

[Section titled “Level 5: Field”]

Results for a single field across all dimensions. Useful for understanding why a specific field is scoring poorly.

Level 6: Dimension

Section titled “Level 6: Dimension”

The deepest level — individual metric results for one field in one dimension. Shows the raw evaluation data.

Sidebar Panel

Section titled “Sidebar Panel”

The right-hand sidebar shows context and actions for the current definition:

- **Header** — Definition name, overall score with grade, scan count, and a refresh button
- **Run Scan** — Trigger a scan manually, with last scan date displayed below
- **Analyze** — **Actions** dropdown menu with **Export Report**, **Create Tasks**, and **Post Chatter** (details)
- **Configure** — **Manage Schedule** with countdown to the next scheduled run (e.g., “Next in 8h 28m”)
- **Navigate** — **Open in Salesforce** (external link to the definition record) and **View Records** (browse scanned records)
- **Quick Fix** — Highlights the worst-performing field with its score (e.g., “Description at 9.6% — worst field”) and a direct **Investigate** link to drill into it

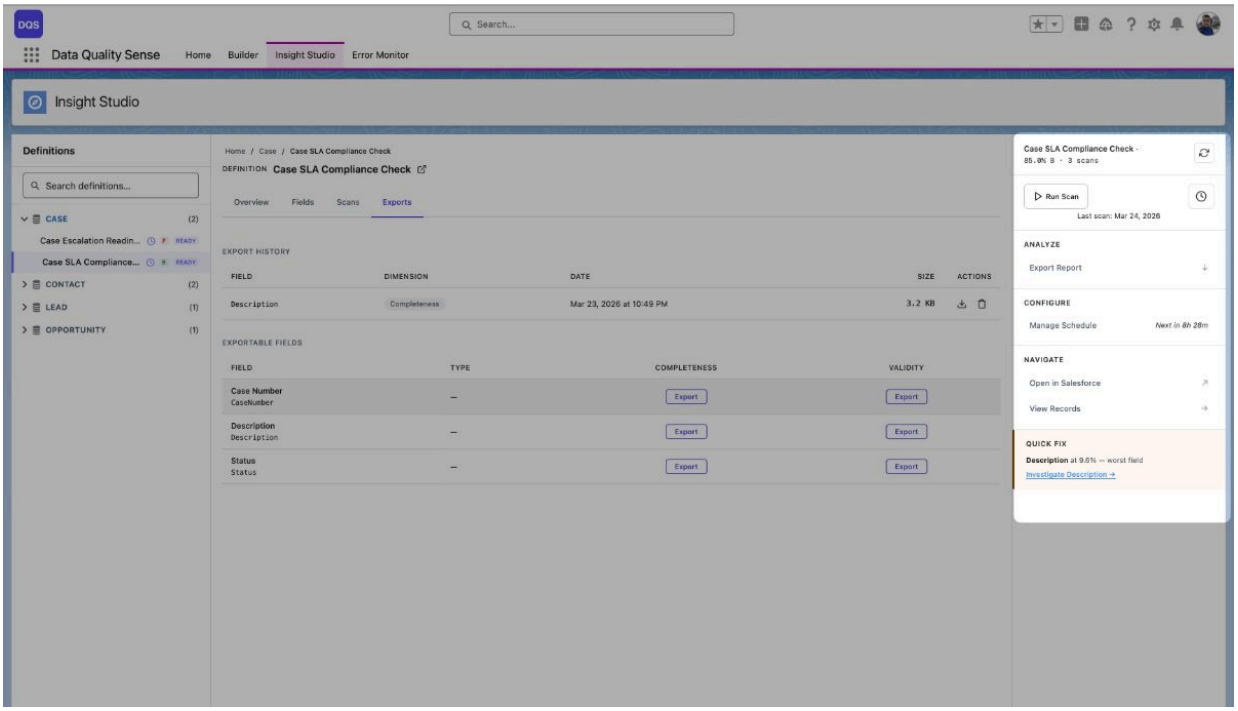

Breadcrumb Navigation

Section titled “Breadcrumb Navigation”

A breadcrumb trail at the top of the stage area shows your current position in the hierarchy. Click any breadcrumb segment to navigate back to that level.

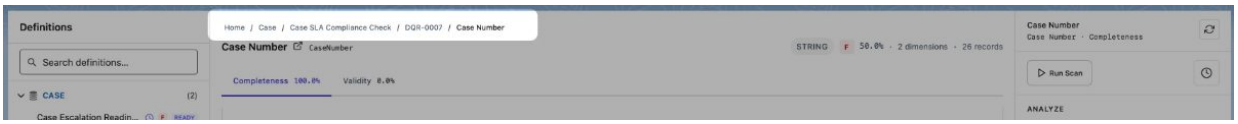

Scores & Trends

Score Display

Section titled “Score Display”

Every level in Insight Studio shows quality scores with visual indicators:

Score Grades

Section titled “Score Grades”

Score Range	Grade	Color
90–100	Excellent	Green
75–89	Good	Light Green
50–74	Fair	Yellow/Orange
25–49	Poor	Orange
0–24	Critical	Red

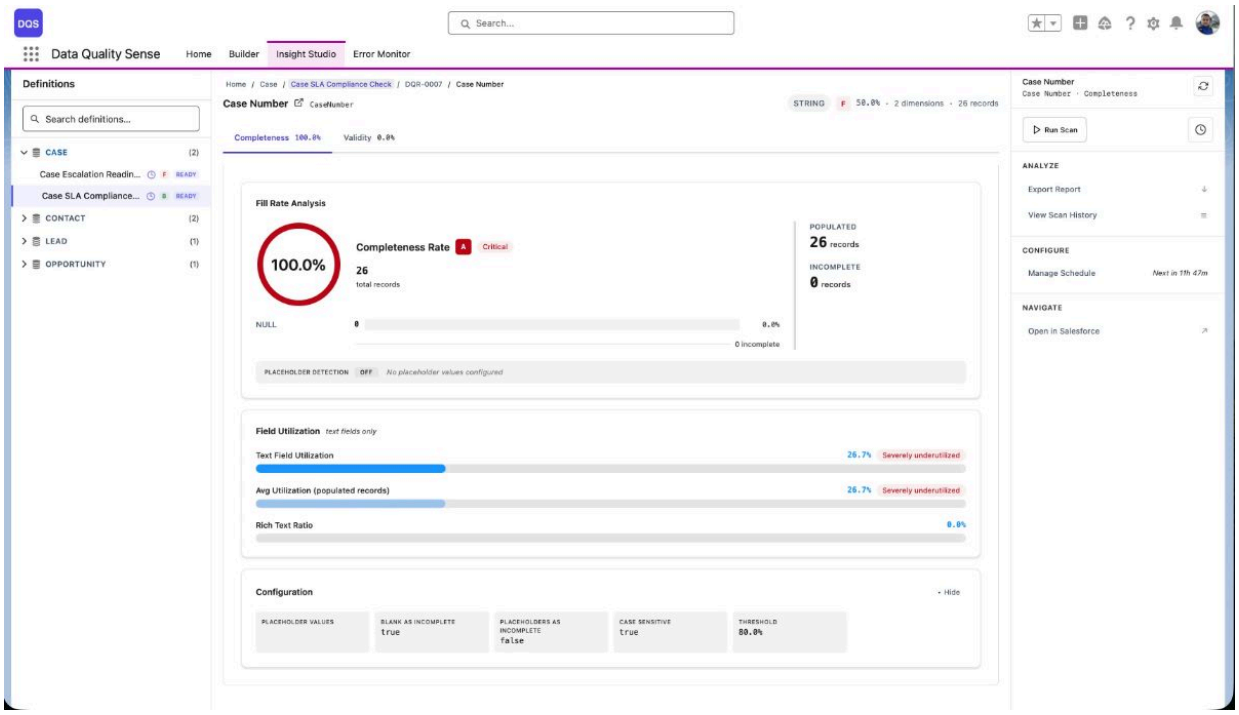

Score Cards

[Section titled “Score Cards”]

The definition dashboard shows **dimension cards** — one for each enabled capability — with:

- Numeric score (0–100)
- Grade indicator (color-coded)
- Change arrow (↑ improvement, ↓ regression, → stable)
- Sparkline showing recent trend

Trend Analysis

[Section titled “Trend Analysis”]

Sparklines

[Section titled “Sparklines”]

Small inline charts next to each score showing the last several scan results. Useful for spotting trends without navigating to detailed views.

Score Trend Chart

[Section titled “Score Trend Chart”]

A larger trend chart available at the definition level showing:

- Score over time for each dimension
- Overall score trend line

- Scan dates on the x-axis

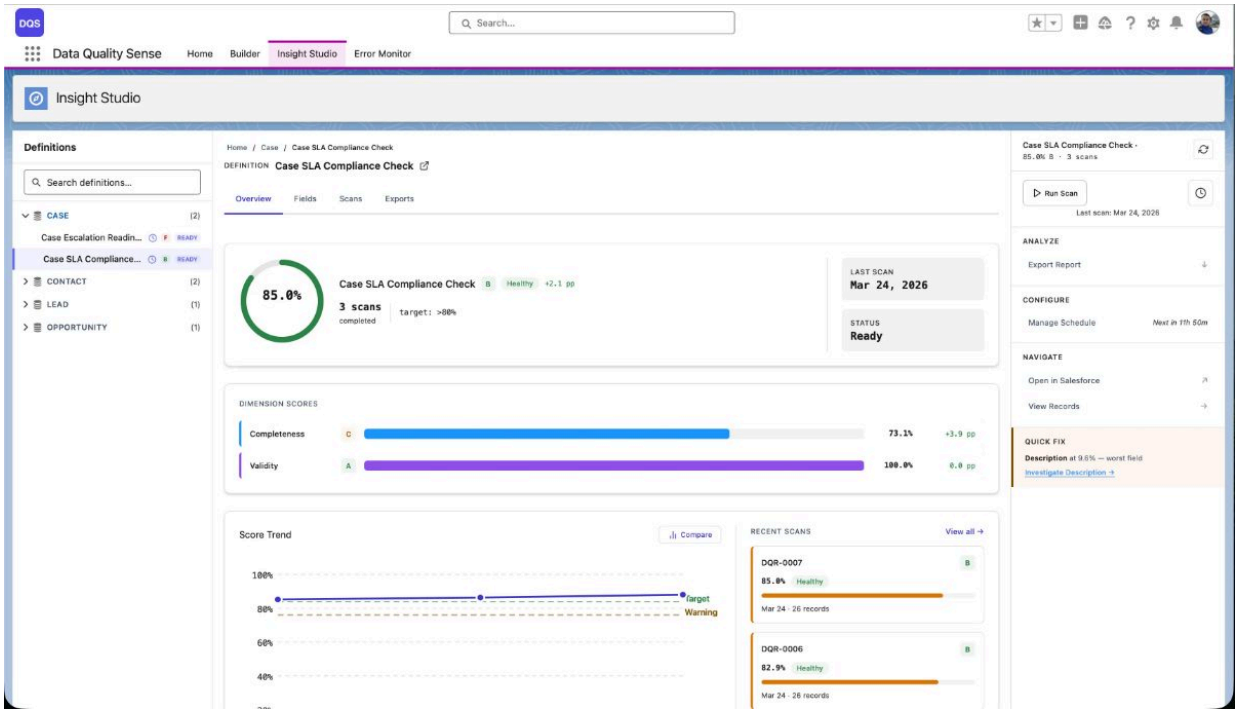

Scan Comparison

[Section titled “Scan Comparison”]

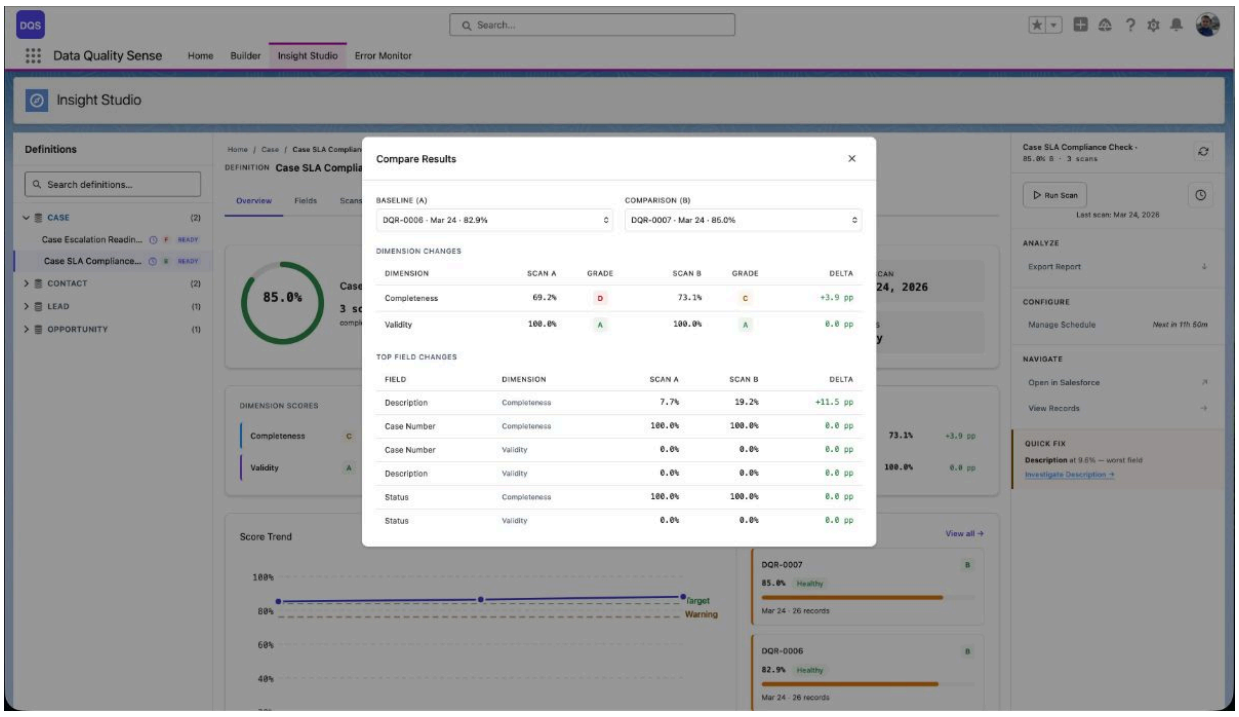

Compare two specific scans to see detailed deltas:

- **Dimension deltas** — Which capabilities improved or regressed
- **Field deltas** — Which fields changed and by how much
- Select any two scans from the scan history

Two-State Views

Section titled “Two-State Views”

Some components adapt based on data availability:

- **No scans yet** — Shows a prompt to trigger the first scan
- **Results available** — Shows the full dashboard with scores and charts

KPI Summary

Section titled “KPI Summary”

The top of the definition view shows key performance indicators:

- **Overall Score** — Weighted average across all dimensions
- **Fields Scanned** — Total fields included
- **Records Evaluated** — Total records in scope
- **Last Scan** — Timestamp of the most recent scan

Field Health

Field Health Matrix

Section titled “Field Health Matrix”

The field health matrix is a **grid view** showing every scanned field against every enabled dimension. Each cell displays the field’s score for that dimension, color-coded by grade.

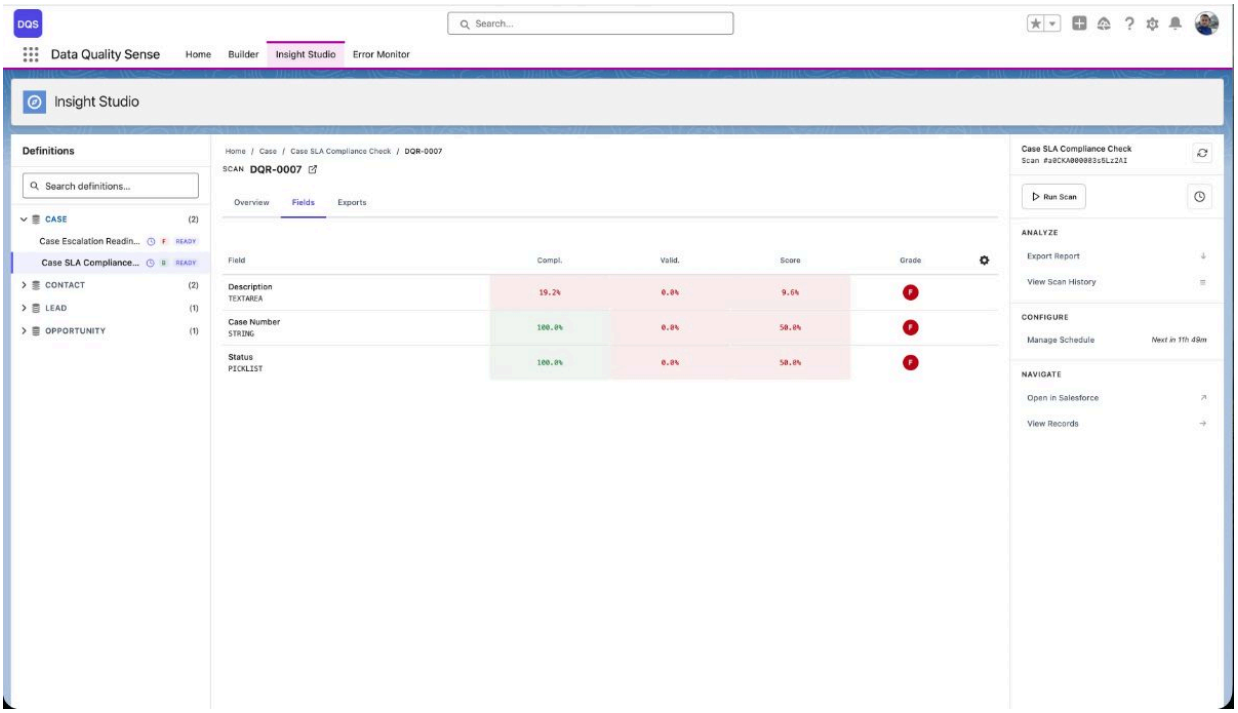

Reading the Matrix

Section titled “Reading the Matrix”

Axis	Content
Rows	Individual fields (sorted by overall score, worst first)
Columns	Quality dimensions (Completeness, Validity, etc.)
Cells	Score (0–100) with color indicator

Color Coding

Section titled “Color Coding”

- **Green cells** — Score ≥ 90 (Excellent)
- **Yellow cells** — Score 50–89 (Fair to Good)
- **Red cells** — Score < 50 (Poor to Critical)
- **Gray cells** — Not applicable (dimension doesn’t apply to this field type)

Using the Matrix

Section titled “Using the Matrix”

Identify Problem Fields

Section titled “Identify Problem Fields”

Scan the matrix for rows with multiple red/yellow cells. These fields have quality issues across multiple dimensions.

Identify Problem Dimensions

Section titled “Identify Problem Dimensions”

Look for columns with many red cells. These dimensions need attention across your dataset.

Drill Down

Section titled “Drill Down”

Click any cell to navigate to the detailed field-dimension view, showing the specific metrics and records that contribute to that score.

Sorting and Filtering

Section titled “Sorting and Filtering”

- **Sort by score** — Show the worst-performing fields first
- **Filter by dimension** — Focus on a specific quality dimension
- **Filter by field type** — Show only Text, Number, or Date fields

---

Actions

Actions Menu

Section titled “Actions Menu”

The **Actions** menu in the Mentor panel provides a unified dropdown for taking action on records that failed quality checks. It replaces the standalone Export button with an extensible menu that currently offers three actions:

Action	What It Does
Export Report	Download violation details as CSV (details)
Create Tasks	Create Salesforce Tasks for record owners to remediate issues
Post Chatter	Post Chatter messages on impacted records to notify stakeholders

All actions share the same scope model — you choose which **fields** and **dimensions** to act on, and the system processes violations for each combination as a separate background job.

Scope Selection

Section titled “Scope Selection”

When you open any action modal, you first select the scope:

1. **Fields** — Choose “All fields” or select specific fields from the definition
2. **Dimensions** — Choose which quality dimensions to include (Completeness, Validity, Uniqueness, etc.)

The system re-evaluates violations at the time of the action, so results reflect the **current state** of your data — not a cached snapshot from the last scan.

Each field-dimension combination runs as a separate batch job. For example, selecting 3 fields and 2 dimensions produces 6 jobs that execute sequentially.

Create Tasks

Section titled “Create Tasks”

Creates Salesforce Task records linked to each impacted record. Tasks are assigned to record owners by default and appear in their standard Salesforce task list.

Task Configuration

Section titled “Task Configuration”

Setting	Default	Description
Subject	`Data Quality: {dimension} – {recordName}`	Task subject line. Supports `{dimension}`, `{recordName}`, and `{fieldName}` placeholders
Description	Describes the quality issue and asks for review	Free-text body of the task
Due Date	7 days from today	When the task should be completed
Priority	Normal	High, Normal, or Low
Assign To	Record Owner	Enter a specific User ID to override. Leave blank to assign to each record’s owner

Duplicate Prevention

Section titled “Duplicate Prevention”

Before creating tasks, the system checks for **existing open tasks** on each record with a matching subject prefix. Records that already have a matching open task are skipped — this prevents duplicate tasks when you run the action multiple times.

The skip count is reported in the completion summary (e.g., “Created 45 tasks. 12 skipped (existing tasks). 0 errors.”).

How It Works

Section titled “How It Works”

1. Open the **Actions** menu in the Mentor panel and select **Create Tasks**
2. Select the fields and dimensions to include
3. Configure the task settings (subject, description, due date, priority, assignee)
4. Click **Create Tasks** — the modal switches to progress mode
5. Each field-dimension job shows its status (pending, running, complete, or failed)
6. When all jobs finish, you receive a **custom notification** with a summary
7. Click **Done** to close the modal

Post Chatter

Section titled “Post Chatter”

Posts a Chatter feed message on each impacted record. Optionally @mentions the record owner to trigger a Salesforce notification.

Chatter Configuration

Section titled “Chatter Configuration”

Setting	Default	Description
Message	Describes the quality issue and dimension	Free-text message body. Supports `{dimension}`, `{recordName}`, and `{fieldName}` placeholders
Mention Record Owner	Checked	When enabled, the post @mentions the record’s owner, triggering a Salesforce notification

Duplicate Prevention

Section titled “Duplicate Prevention”

Chatter uses a **24-hour dedup window** — if the current user already posted a matching quality message on a record within the last 24 hours, that record is skipped. This prevents flooding Chatter feeds when running the action repeatedly.

How It Works

Section titled “How It Works”

1. Open the **Actions** menu and select **Post Chatter**
2. Select the fields and dimensions to include
3. Edit the message template and choose whether to @mention owners
4. Click **Post Messages** — the modal switches to progress mode
5. When complete, the summary shows posted/skipped/error counts

6. Click **Done** to close the modal

ⓘNote

Post Chatter requires Chatter to be enabled in your Salesforce org. If Chatter is not available, the action will report an error with a clear message.

Processing Details

🔗Section titled “Processing Details”

All actions run as **Apex batch jobs** in the background. Key details:

- **Batch size:** 500 records per chunk (smaller than export to allow for DML operations per chunk)
- **Notifications:** A Salesforce custom notification is sent when the batch completes, showing created/skipped/error counts
- **Partial success:** If some records fail (e.g., validation rules, field-level security), successfully processed records are preserved. The error count and details are included in the summary
- **Batch queue:** Salesforce allows up to 5 concurrent batch jobs per org. If the queue is full, you will receive feedback before the action starts

Use Cases

🔗Section titled “Use Cases”

- **Task creation for data stewards** — Automatically assign cleanup tasks to record owners when quality issues are detected
- **Stakeholder notification via Chatter** — Alert record owners about data quality problems directly on their records
- **Remediation workflows** — Combine with Salesforce automation (Flow, Process Builder) to route tasks to specific teams
- **Audit trails** — Created tasks and Chatter posts serve as a record of quality issues and when they were flagged

Exports

CSV Export

🔗Section titled “CSV Export”

Insight Studio supports **CSV export** of violation details for all dimensions. This lets you take action on data quality issues outside of Salesforce.

What Gets Exported

Section titled “What Gets Exported”

Each export contains the specific records that failed quality checks for a given dimension:

Dimension	Export Contains
Completeness	Records with blank/null fields
Validity	Records with invalid values
Uniqueness	Records with duplicate values
Timeliness	Records with stale dates
Consistency	Records with contradicting fields
PII Detection	Records with detected PII

Export Columns

Section titled “Export Columns”

Each CSV includes:

- **Record ID** — Salesforce record identifier
- **Record Name** — Display name of the record
- **Field** — The field that failed the check
- **Current Value** — The value that triggered the violation
- **Violation Type** — Specific type of quality issue
- **Score** — The field’s score for this dimension

How to Export

Section titled “How to Export”

1. Navigate to the dimension or field view in Insight Studio
2. Click the **Export** button
3. The export runs as a background process — you receive a **custom notification** in the bell icon when it completes. Each notification shows the dimension, field, and number of violations exported (e.g., “Validity — Status: 0 violations exported”)
4. Click the notification to navigate to the **Field Result** page in Salesforce, where the CSV is attached in the **Files** panel

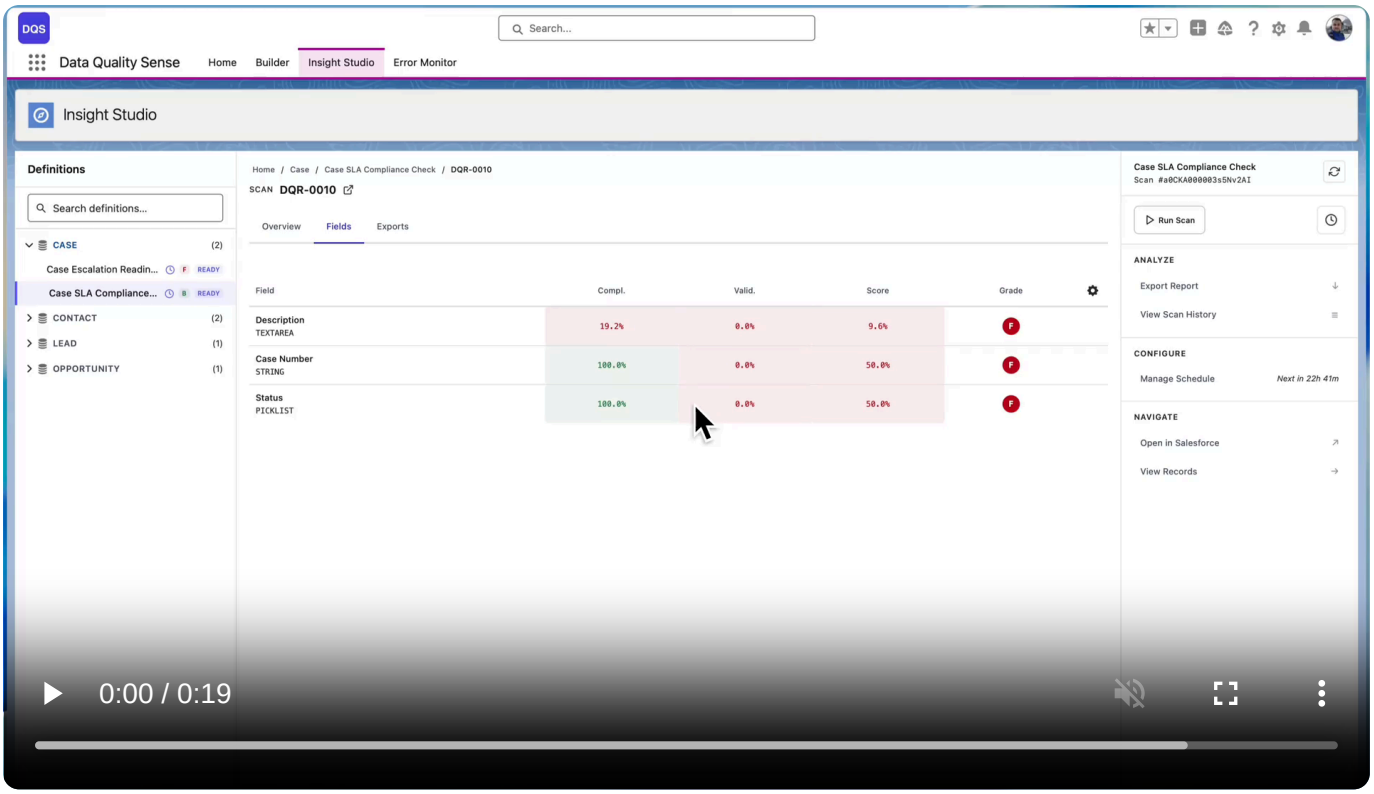

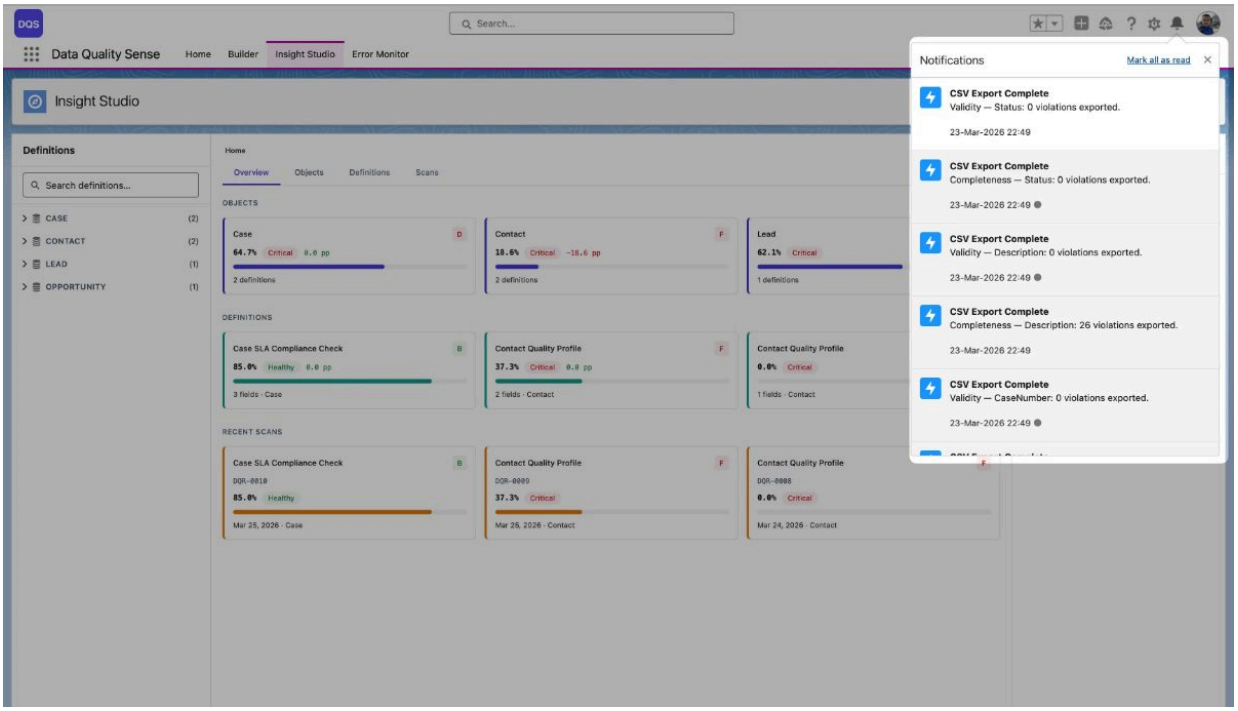

The Field Result page shows the full context — field identity, dimension, scoring details (score, threshold, records empty/populated), and an **AI Insight** section. The exported CSV file is available for download in the **Files** panel on the right.

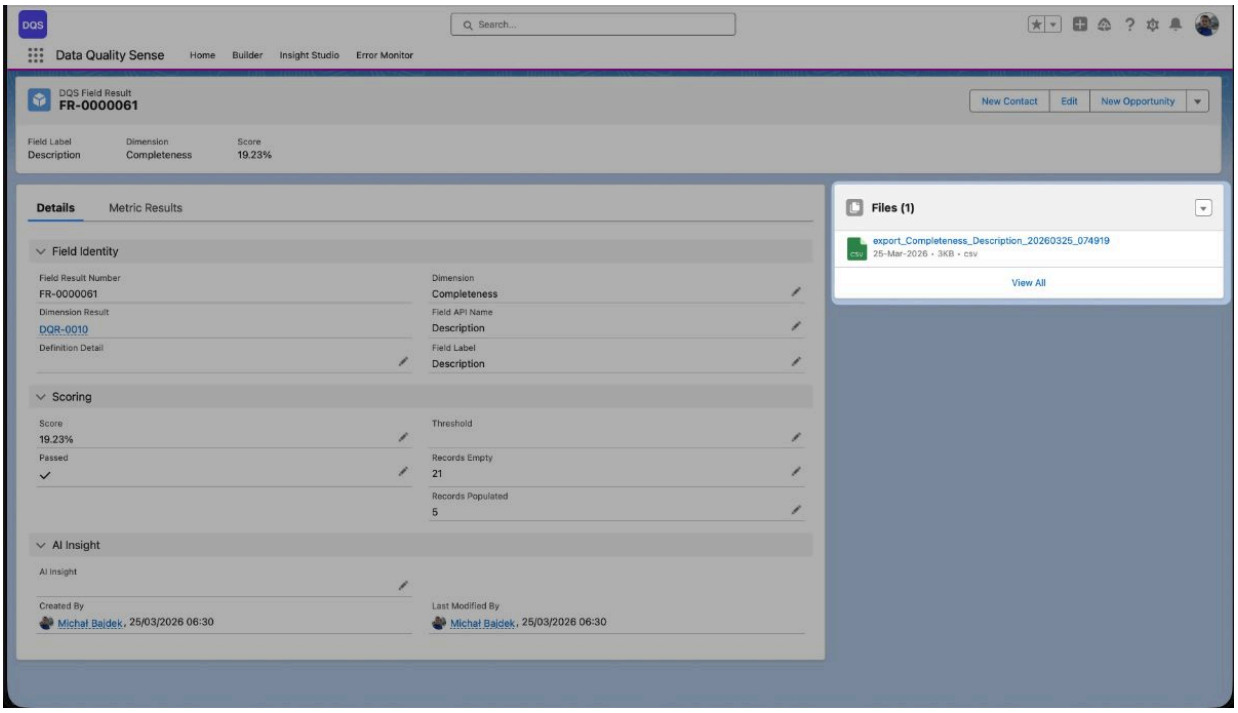

Below is an example of the exported CSV file contents:

export_Completeness_Description_20260323_214940

RecordUrl	FieldName	Value	ViolationType	Details
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTbAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTcAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTdAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTeAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTfAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTgAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwThAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTiAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTjAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTkAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTlAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTmAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTnAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwToAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTpAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTqAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTrAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTsAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTtAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTuAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTvAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTwAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTxAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTyAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwTzAAK/view	Description		NULL	Field value is null
https://mb-dev-env-dev-ed.my.salesforce.com/lightning/r/Case/5001r00001xMwU0AAK/view	Description		NULL	Field value is null

Use Cases

Section titled “Use Cases”

- Share data quality reports with stakeholders who don’t have Salesforce access
- Import violation lists into data cleansing tools
- Create action items in project management tools
- Compliance documentation and audit trails

> **Processing**

Processing Overview

How Scans Work

Section titled “How Scans Work”

The Processing engine is the **execution layer** of Data Quality Sense. It reads scan definitions created in the Builder, queries live Salesforce data, and produces quality scores visible in Insight Studio.

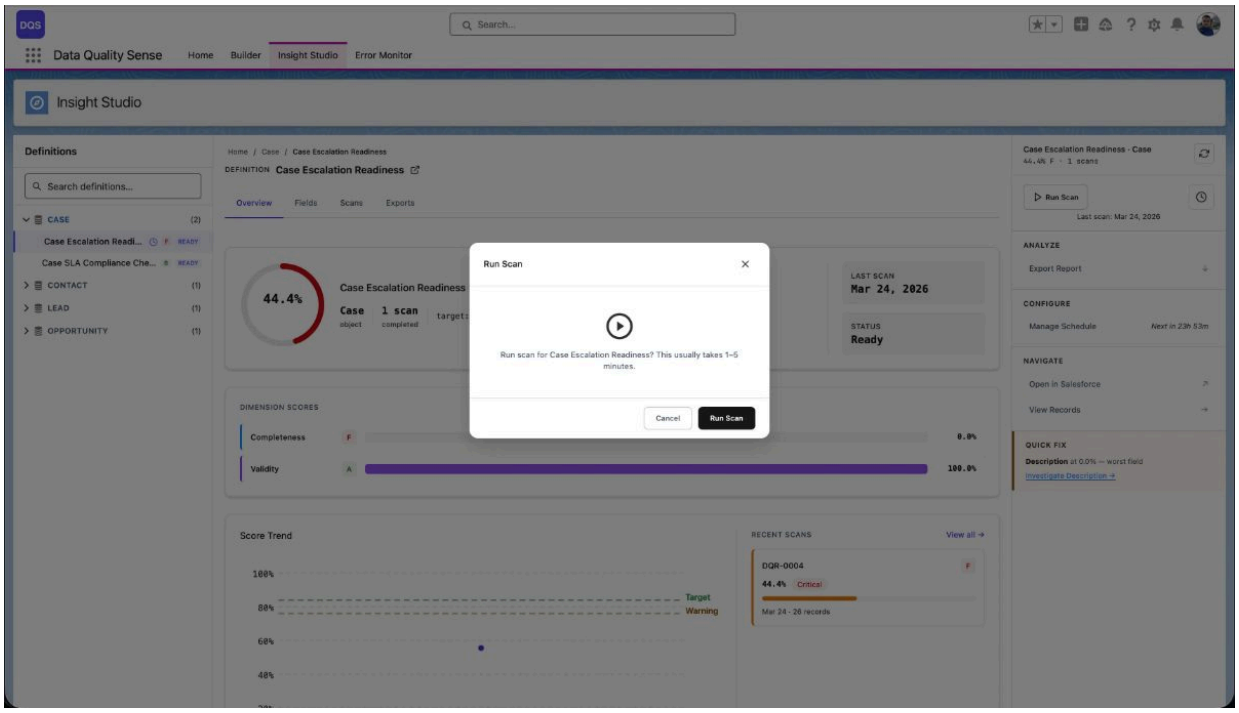

Triggering a Scan

Section titled “Triggering a Scan”

Scans can be triggered in three ways:

1. **Manual** — Click “Run Scan” in Insight Studio
2. **Scheduled** — Via CRON-based scheduling
3. **Programmatic** — Via Apex (for custom integrations)

Learn More

Section titled “Learn More”

- Scheduling
- Data Retention
- Error Management

Scheduling

Scan Scheduling

Section titled “Scan Scheduling”

DQS supports **automated recurring scans** via CRON-based scheduling. Once configured, scans run automatically without manual intervention.

Creating a Schedule

Section titled “Creating a Schedule”

1. Navigate to the definition in **Insight Studio**
2. Open the **Scan Schedules** modal from the sidebar
3. Click **Create Schedule**

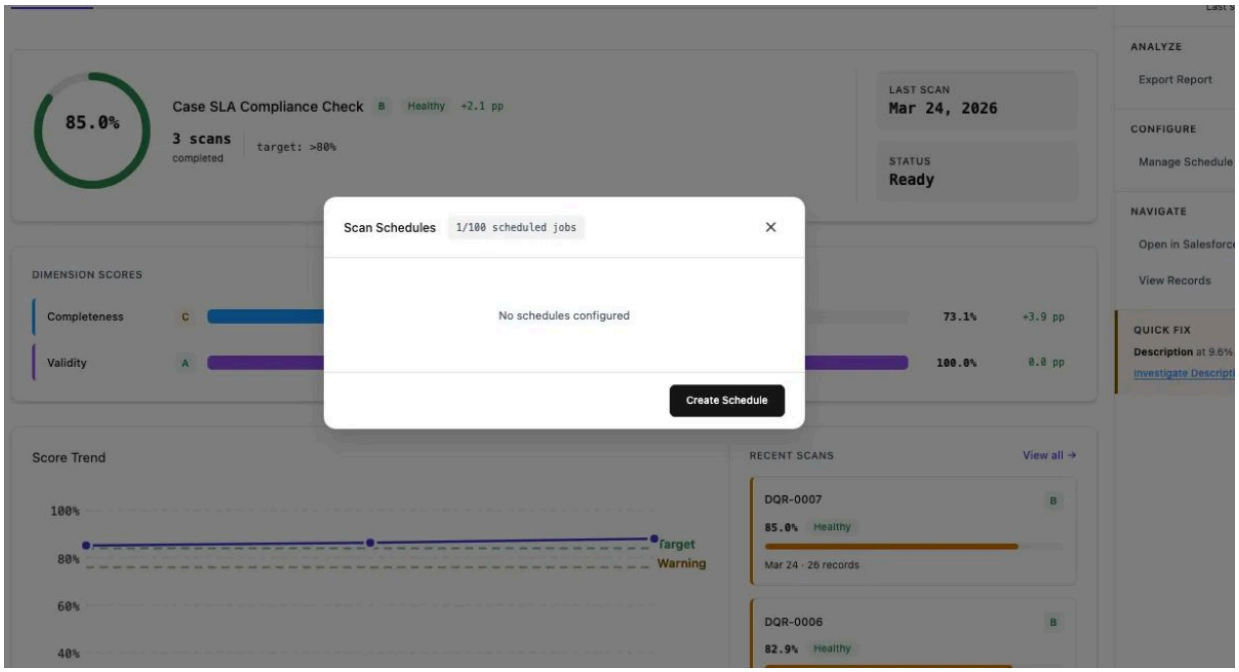

4. Configure the schedule — set the frequency, time, and optionally a name:

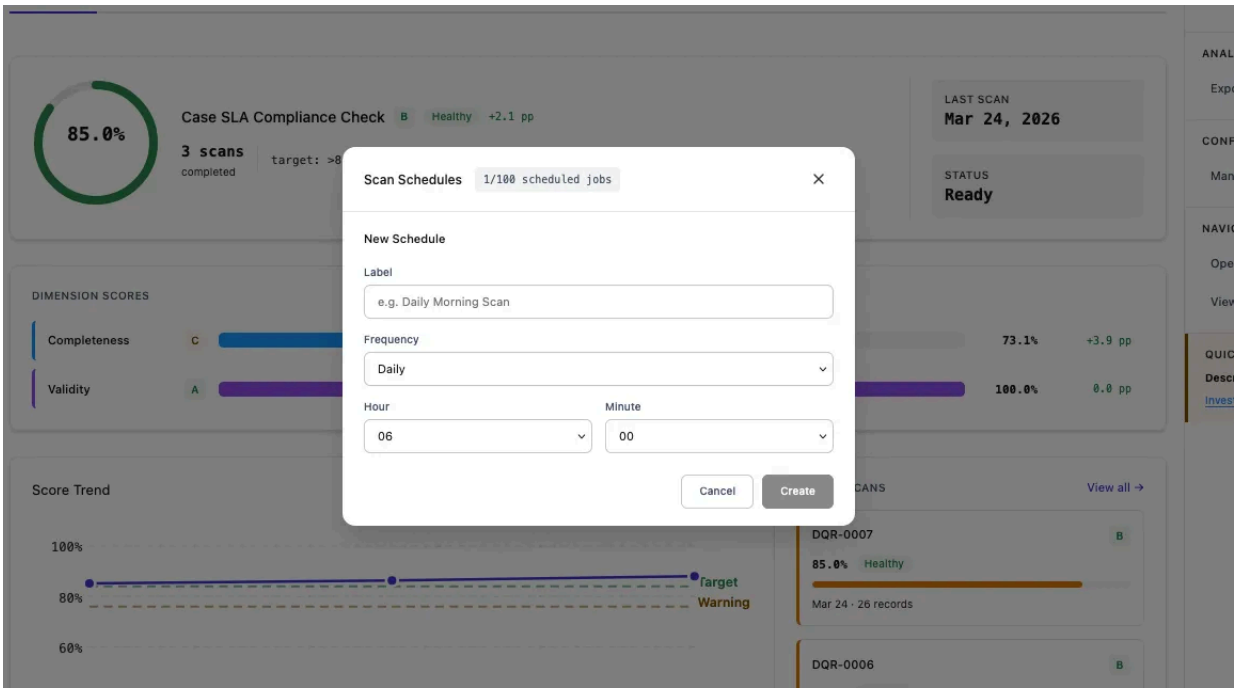

5. Review and **Save** the schedule:

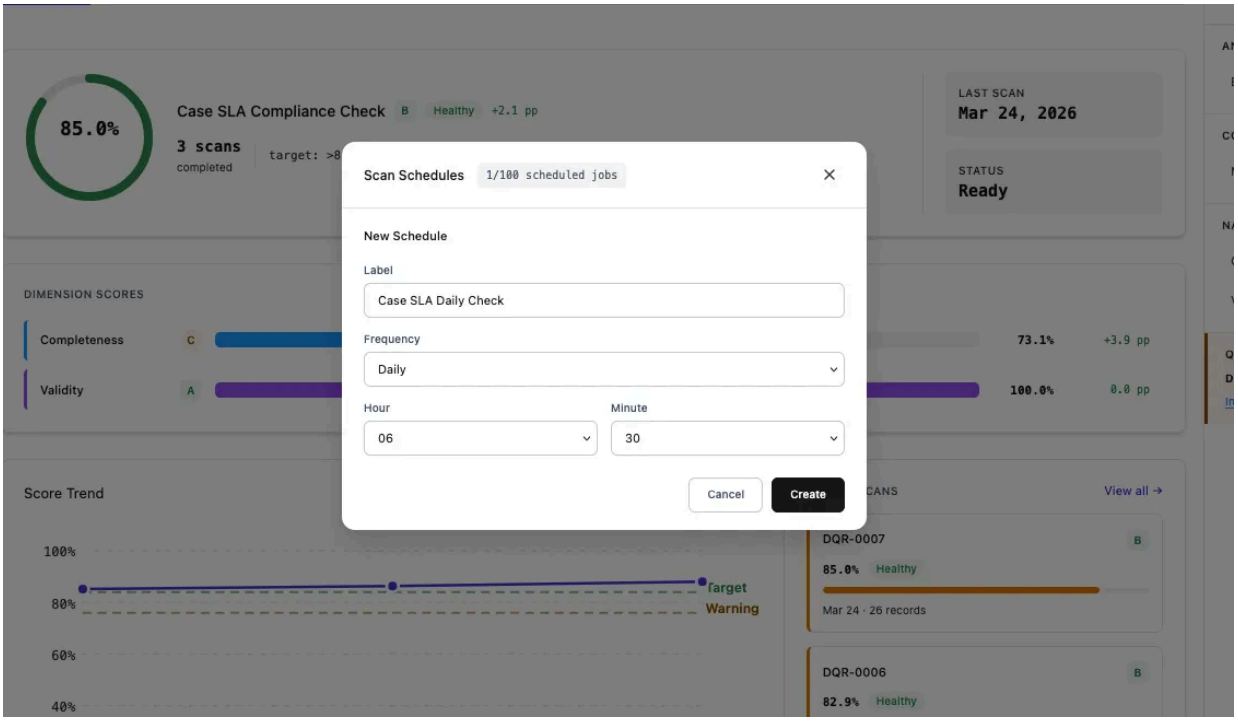

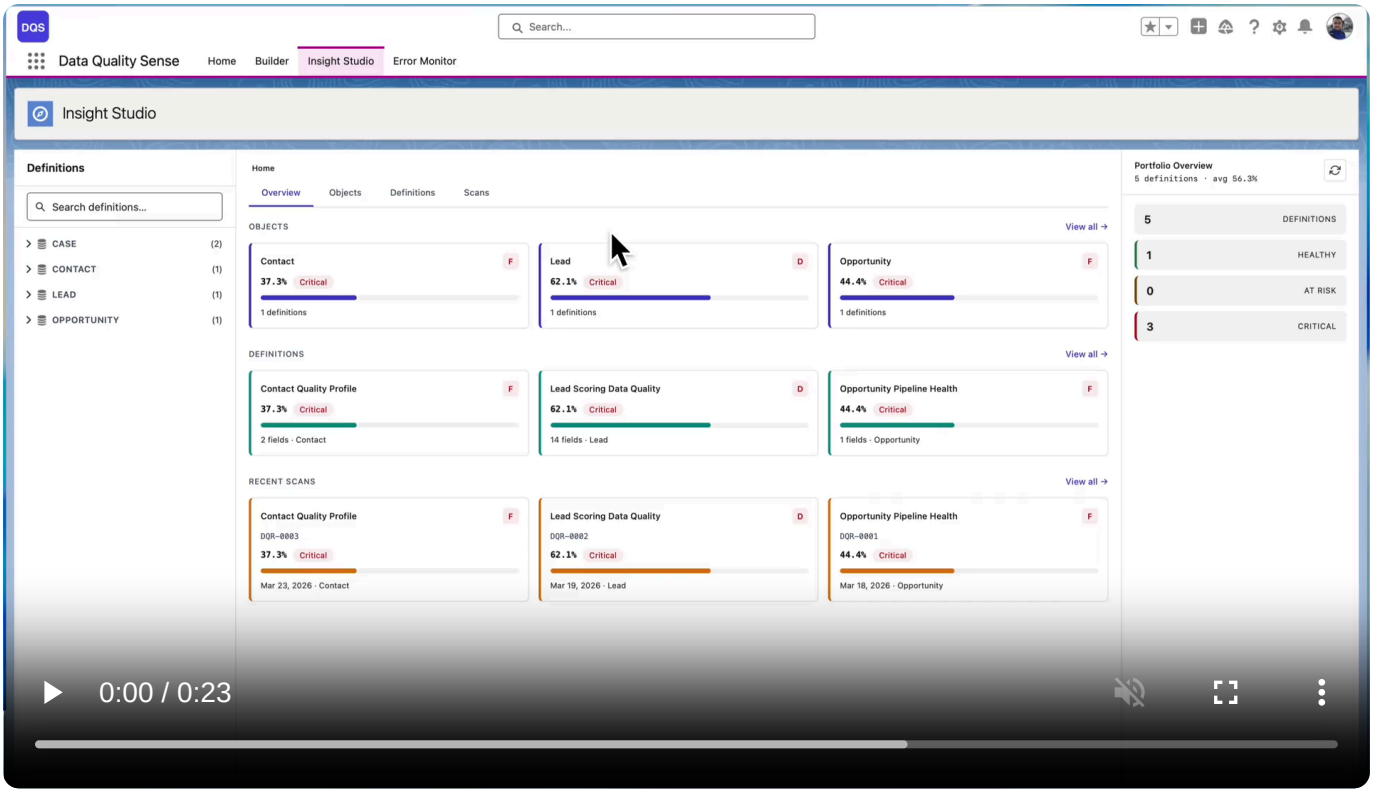

Schedule Settings

Section titled “Schedule Settings”

Setting	Description	Example
Name	Display name for the schedule	Case SLA Daily Check
Frequency	How often the scan runs	Daily, Weekly, Monthly
Time	What time of day to run	06:30
Day of Week	For weekly schedules	Monday
Day of Month	For monthly schedules	1st

CRON Expressions

Section titled “CRON Expressions”

Under the hood, schedules use Salesforce CRON expressions. DQS provides a **user-friendly UI** that generates the CRON expression for you, but advanced users can also set custom expressions.

Common Schedules

Section titled “Common Schedules”

Schedule	CRON Expression
Daily at 2 AM	`0 0 2 * * ?`
Weekly on Monday at 6 AM	`0 0 6 ? * MON`

Schedule	CRON Expression
Monthly on 1st at midnight	`0 0 0 1 * ?`
Every weekday at 5 AM	`0 0 5 ? * MON-FRI`

Schedule Management

Section titled “Schedule Management”

Viewing Schedules

Section titled “Viewing Schedules”

The Scan Schedules modal shows all configured schedules with their status, frequency, and next run time. The counter shows how many scheduled jobs are used out of the Salesforce limit (e.g., 2/100).

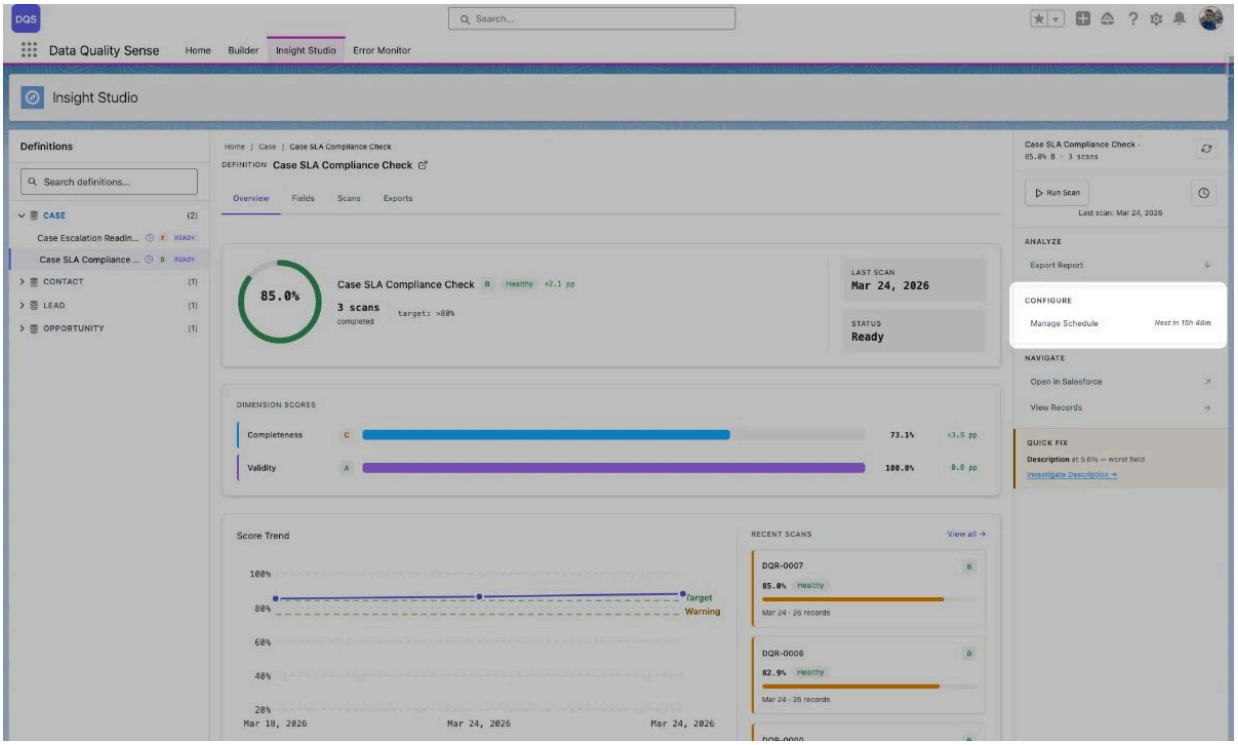

Editing Schedules

Section titled “Editing Schedules”

Click the edit icon to modify schedule settings. The existing schedule is replaced with the new configuration.

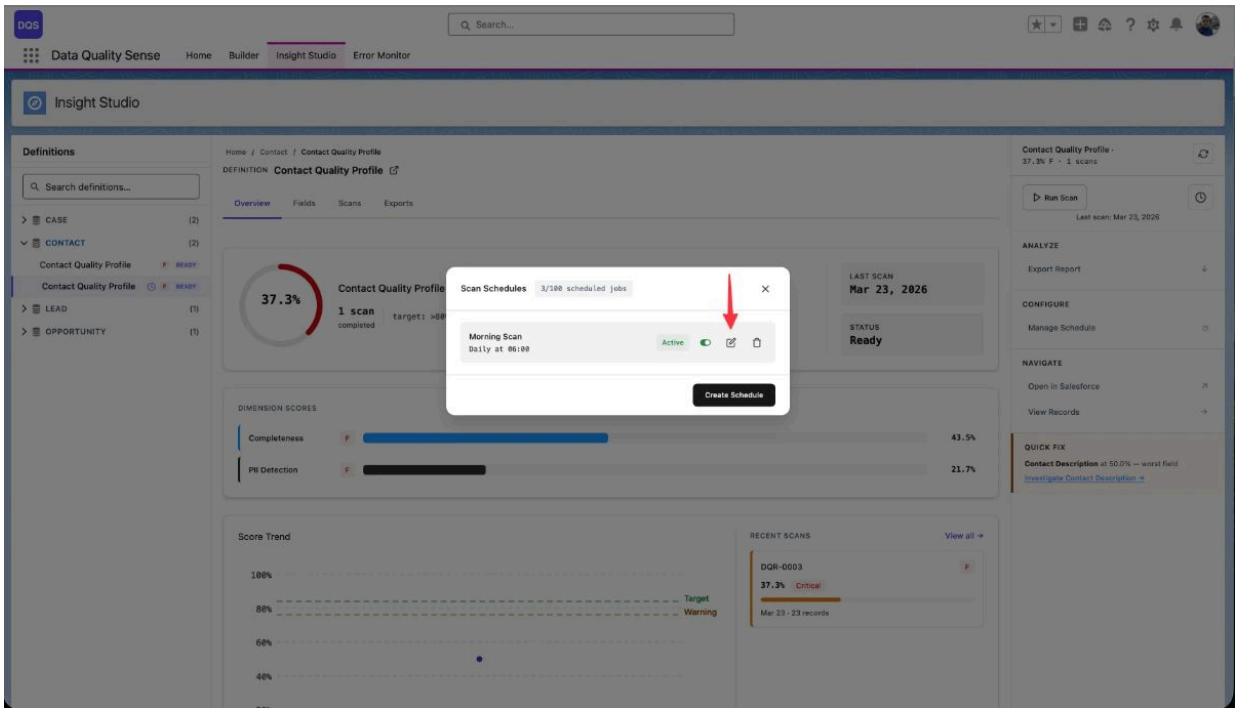

Activating / Deactivating

[Section titled “Activating / Deactivating”]

Use the **Active** toggle to temporarily pause a schedule without deleting it.

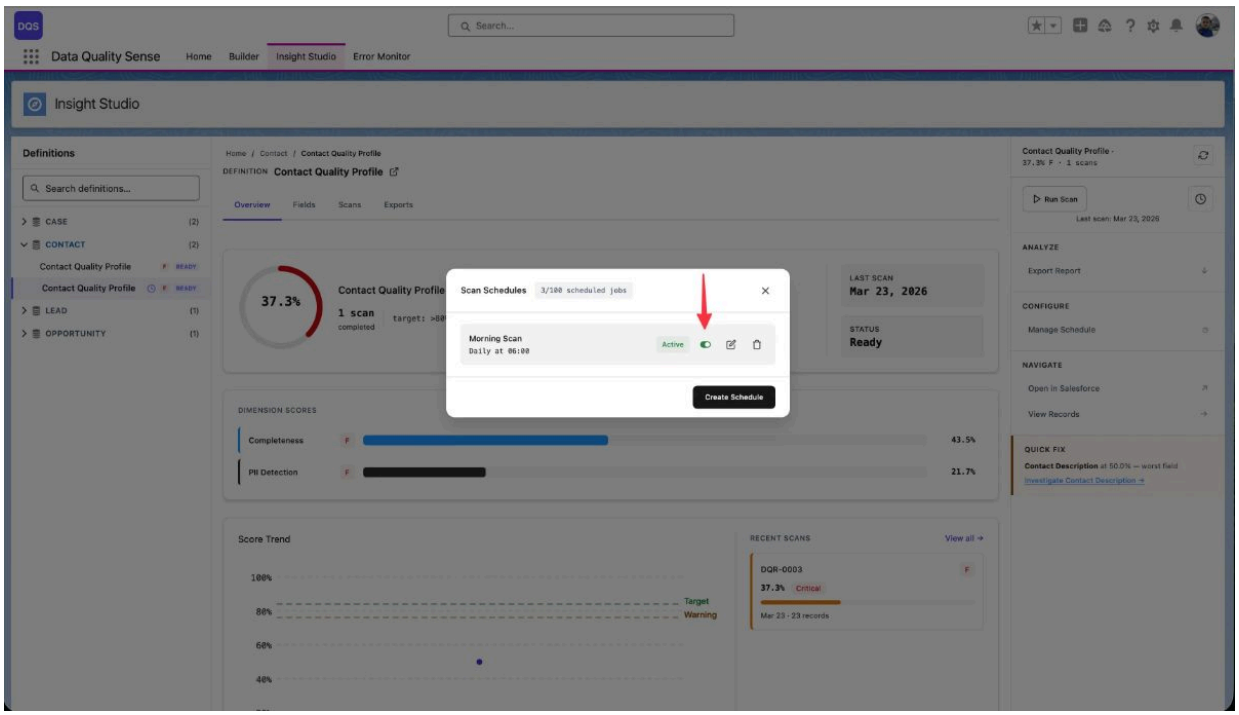

Deleting Schedules

[Section titled “Deleting Schedules”]

Click the delete icon to remove a schedule. Manual scans remain available.

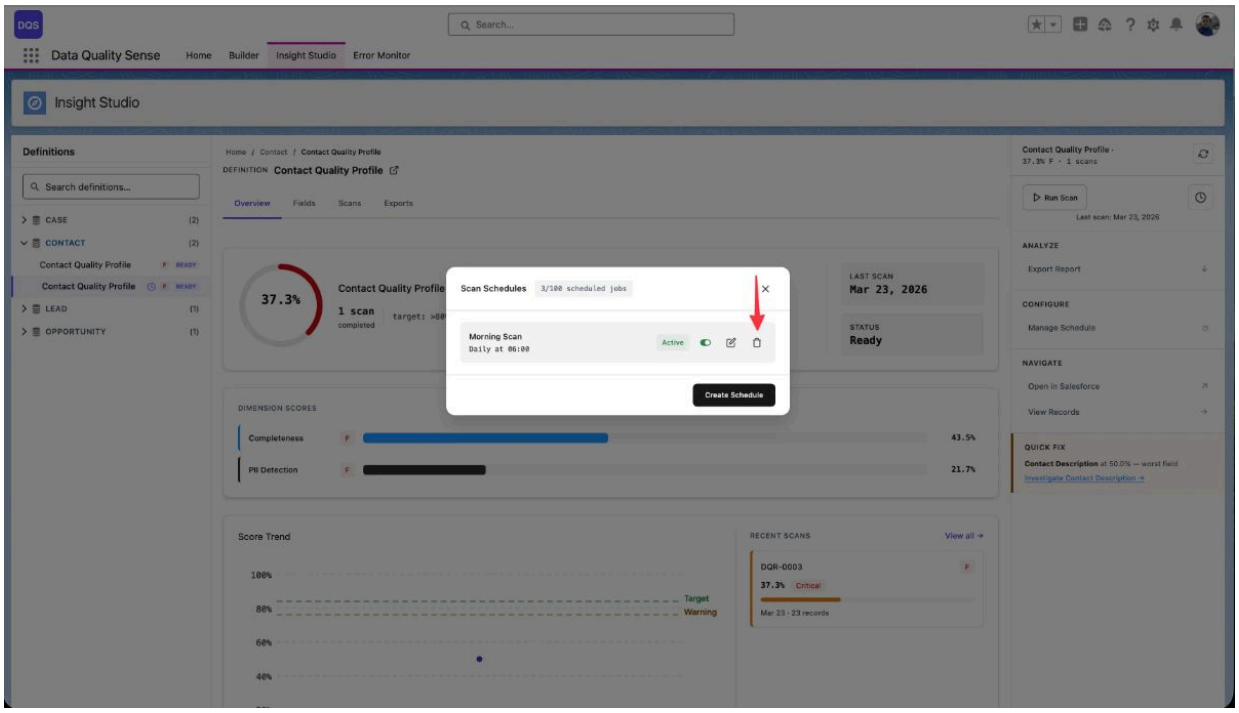

Note

Schedule times use your Salesforce org's timezone. The countdown display shows the correct local time.

Audit Trail

Section titled "Audit Trail"

Every scan records who triggered it:

- **Manual scans** — Records the user who clicked "Run Scan"
- **Scheduled scans** — Records "Scheduled" as the trigger source
- **Timestamp** — When the scan started and completed

Considerations

Section titled "Considerations"

- Only **Active** definitions can be scheduled
- Running scans consume Salesforce governor limits (batch Apex)
- Schedule scans during off-peak hours to minimize impact
- Multiple definitions can be scheduled at different times

Data Retention

Retention Policies

Section titled “Retention Policies”

DQS includes **automated data purging** to prevent unbounded growth of scan results and error logs. Retention policies are configured via Custom Metadata Types and can be adjusted by administrators.

Configuration

Section titled “Configuration”

All retention settings are stored in `DQS_Configuration__mdt` (Category: “Retention”):

Setting	Default	Description
Error Log Retention	7 days	Days before error logs are deleted
Scan Result Retention	30 days	Days before dimension results are purged
Purge Batch Size	2,000	Records processed per batch chunk
Purge CRON Expression	`0 0 2 * * ?`	When the purge job runs (default: daily at 2 AM)

Tip

These settings are Subscriber Controlled — administrators can override default values after package installation.

How Purging Works

Section titled “How Purging Works”

The purge process runs as a **chained batch job**:

```
DQS_DataPurgeScheduler (CRON trigger)
  └── DQS_ErrorLogPurgeBatch
       (deletes error logs where Expires_At <= NOW)
       └── DQS_ResultPurgeBatch
            (deletes dimension results older than retention window)
            └── Cascade: Field Results + Metric Results
                 (deleted automatically via master-detail relationship)
```

Cascade Deletion

Section titled “Cascade Deletion”

When a `DQS_Dimension_Result__c` record is deleted:

- All child `DQS_Field_Result__c` records are automatically deleted
- All grandchild `DQS_Metric_Result__c` records are automatically deleted

This happens via Salesforce's master-detail cascade — no additional batch processing is needed.

Adjusting Retention

Section titled "Adjusting Retention"

To change retention periods, edit the values directly from the **Configuration** panel on the DQS Home page:

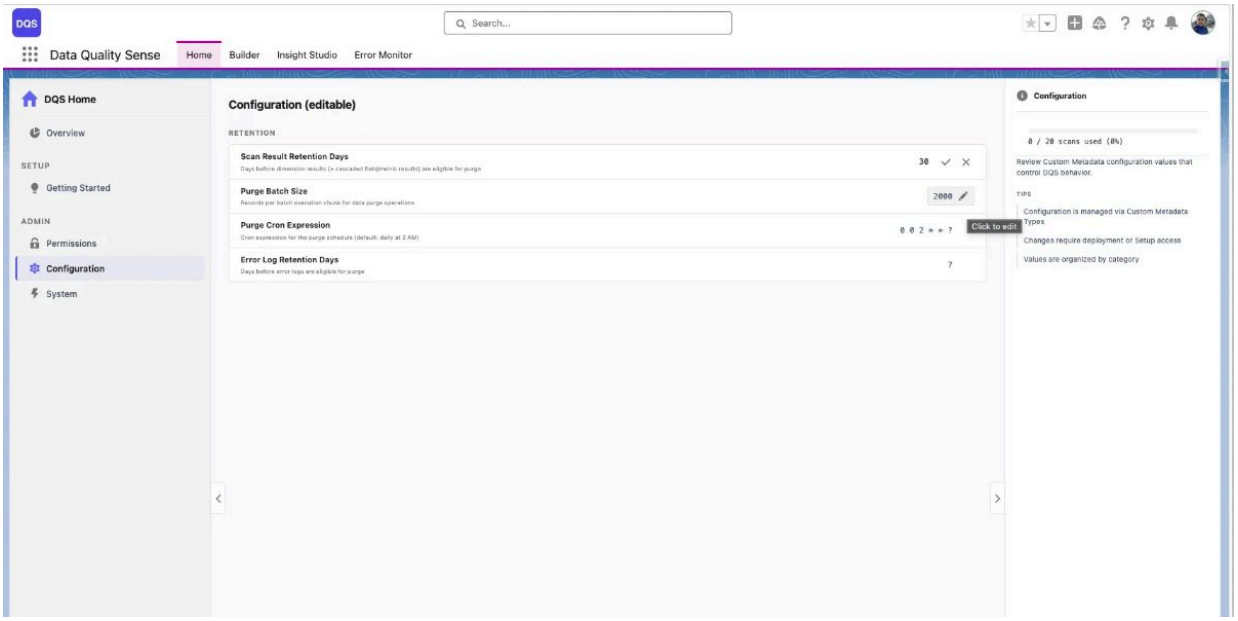

Alternatively, you can update them via Salesforce Setup:

1. Navigate to **Setup → Custom Metadata Types → DQS Configuration**
2. Find the relevant record (e.g., `Scan_Result_Retention_Days`)
3. Edit the value
4. Changes take effect on the next purge run

Monitoring

Section titled "Monitoring"

- Check the **Error Management Console** for purge job errors
- Review `DQS_DataPurgeScheduler` in **Setup → Scheduled Jobs** to verify the purge is running
- Use Salesforce reports to monitor result volume over time

Error Management

Error Management Console (EMC)

Section titled “Error Management Console (EMC)”

The **Error Monitor** tab in the DQS application provides a dedicated interface for monitoring and resolving errors that occur during scan processing. It gives full visibility into batch failures, strategy errors, and platform issues — all in one place.

Layout

Section titled “Layout”

The Error Management Console uses a 3-zone layout:

- **Filters sidebar** (left) — Filter errors by Error Type, Source Class, Error Message, time range (From/To), and IP address. A **Breakdowns** section groups errors by type, source, or other criteria, showing counts next to each category for quick triage.
- **Error log table** (center) — Sortable table listing all errors with summary columns. Displays total error counts across three time windows: Last 24 Hours, Last 7 Days, and Expiring Soon. Click any row to view full details.
- **Actions panel** (right) — Refresh the log, create a downloadable CSV export, delete selected entries, and configure the retention period.

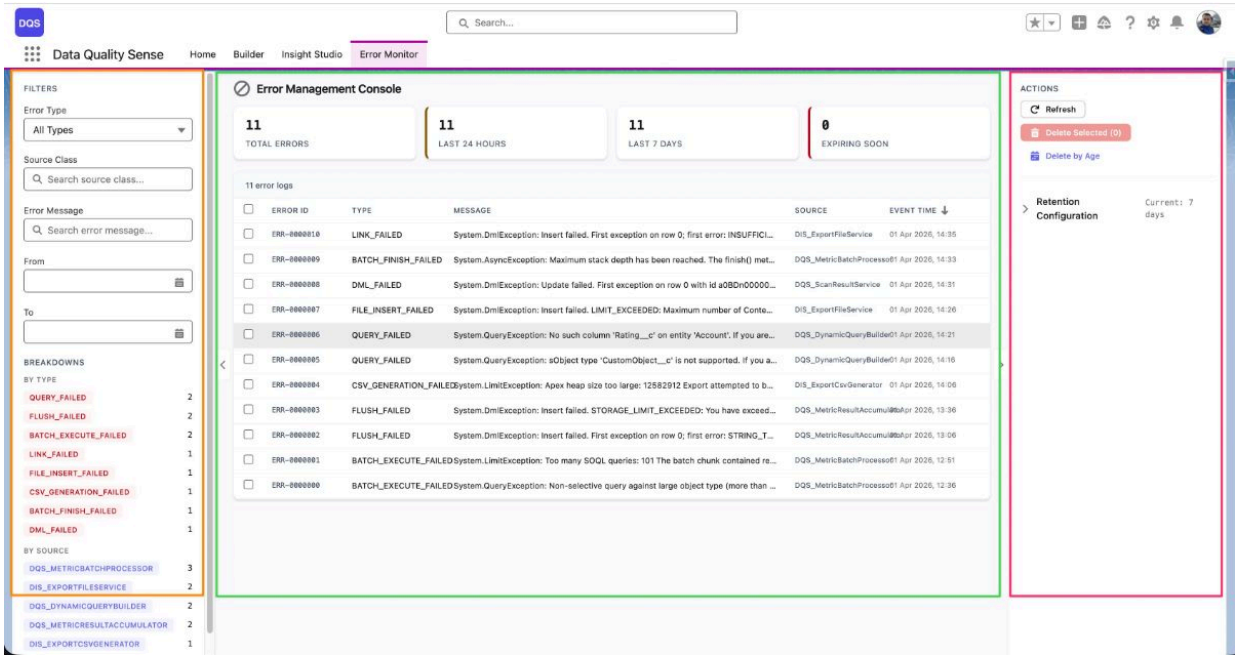

Error Log Table

Section titled “Error Log Table”

Each error row in the table shows the following columns:

Column	Description
Error ID	Unique identifier for the error log entry
Type	Error category (e.g., `DML_FAILED`, `QUERY_FAILED`, `FIELD_INSERT_FAILED`, `BATCH_EXECUTE_FAILED`)
Message	Short error description (truncated — click the row to see the full message)
Source	The Apex class that generated the error (e.g., `DQS_ExportFileService`, `DQS_DynamicQueryBuilder01`)
Event Time	Timestamp of when the error occurred

The top of the table shows aggregate counters — **Total Errors**, **Last 24 Hours**, **Last 7 Days**, and **Expiring Soon** — giving you an at-a-glance health overview.

Row Detail Modal

Section titled “Row Detail Modal”

Click any error row to open a detail modal with the full context:

- **Error ID** and **Type** (highlighted in red for quick identification)
- **Source** — the originating Apex class
- **uniqueIdentifier_01** — the related record or batch identifier
- **Insert Type** — whether the operation was an insert, update, or other DML type
- **Event Time** and **First Event Time** — timestamps for the error occurrence
- **Message** — full error message including the Salesforce exception type (e.g., `System.AsyncException`, `System.DmlException`)
- **Stack Trace** — complete Apex stack trace for debugging

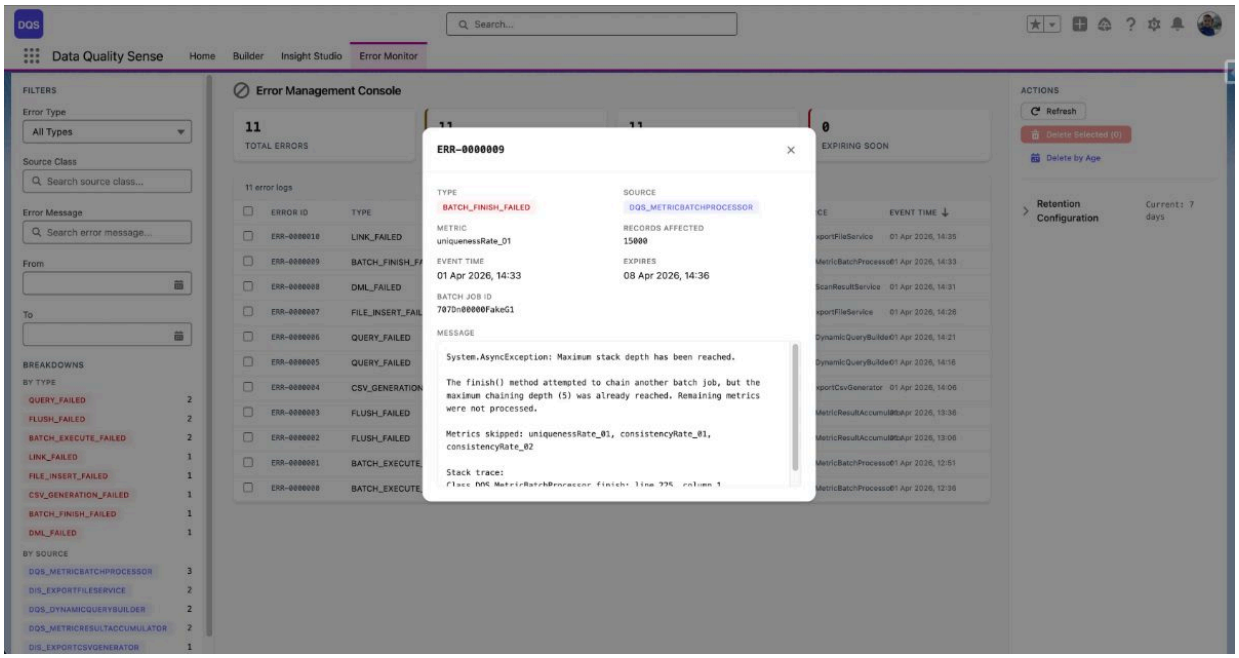

Error Sources

Section titled “Error Sources”

Source	Examples
Batch Processing	Governor limit exceeded, query timeout
Dimension Strategy	Invalid field access, null pointer in strategy logic
Platform Events	Event publish failure
Scheduling	CRON expression issues, permission errors

Platform Event Integration

Section titled “Platform Event Integration”

DQS uses `DQS_Processing_Error__e` platform events to surface errors in real time. When an error occurs during batch processing:

1. The error is caught and logged
2. A platform event is published
3. The EMC receives the event and displays it

This replaces silent `catch` blocks with visible error reporting.

Error Retention

Section titled “Error Retention”

Error logs are automatically purged based on the configured retention period. The default retention is **7 days**.

To change the retention period, use the **Retention Configuration** section in the Actions panel on the right side of the console. Enter the desired number of days in the **Error Log Retention (days)** field and click **Save**. Error logs older than the specified period will be automatically purged. Changes are deployed via the Metadata API and may take a moment to take effect.

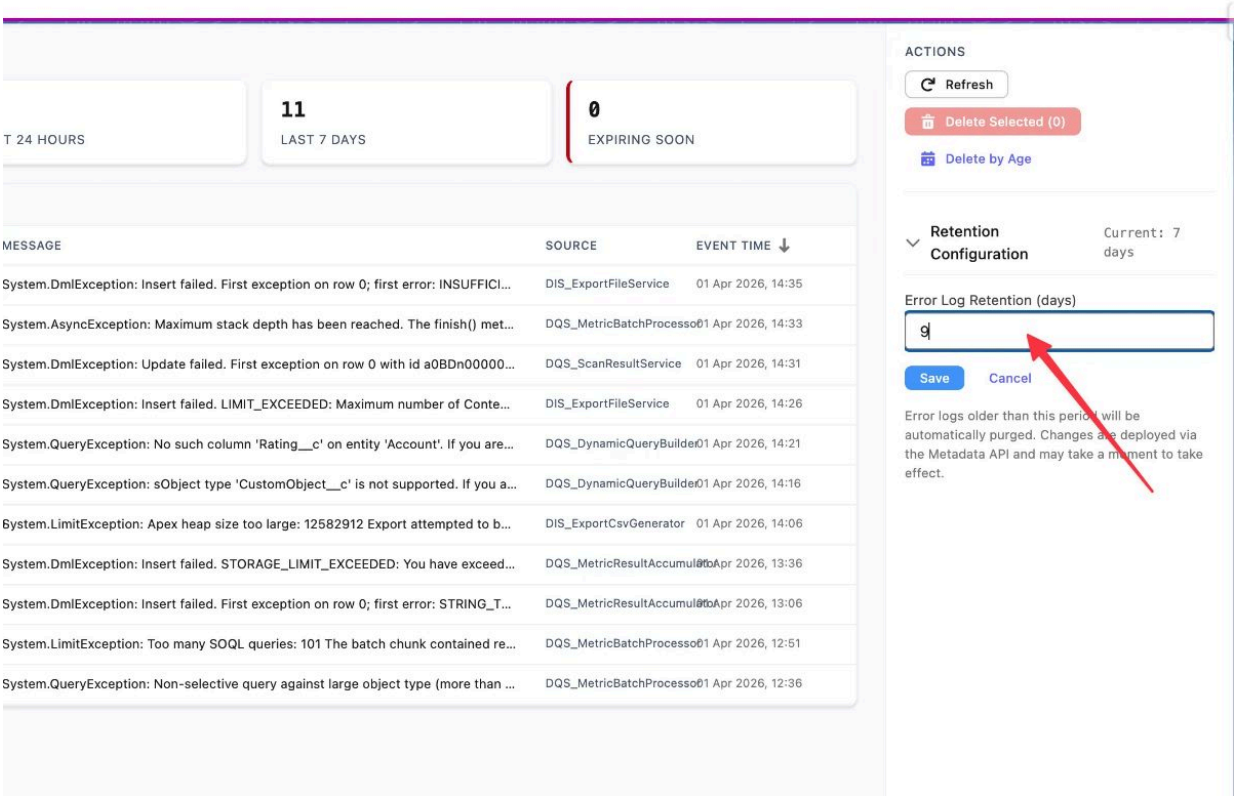

Best Practices

Section titled “Best Practices”

- Check the EMC regularly after setting up new scan definitions
- Review errors after changing capability configurations
- Use error patterns to identify systematic issues (e.g., field access permissions)
- Use the **Expiring Soon** counter to review errors before they are automatically purged
- Export error logs to CSV before they expire if you need to keep them for longer analysis
- Contact support if platform errors persist across multiple scans

> **Reference**

Data Model

Custom Objects

Section titled “Custom Objects”

DQS uses a set of custom objects to store definitions, results, and scheduling data. All objects use the `dataqualitysens` namespace prefix.

Definition Objects

Section titled “Definition Objects”

Object	Purpose	Key Fields
`DQS_Definition__c`	Scan configuration	Name, Object API Name, Status, Description
`DQS_Definition_Detail__c`	Field-level config	Definition (lookup), Field API Name, Overrides

Result Objects

Section titled “Result Objects”

Object	Purpose	Relationship
`DQS_Dimension_Result__c`	Per-dimension scan result	Definition (lookup)
`DQS_Field_Result__c`	Per-field result within a dimension	Dimension Result (master-detail)
`DQS_Metric_Result__c`	Detailed metric per field	Field Result (master-detail)

Scheduling Objects

Section titled “Scheduling Objects”

Object	Purpose
`DQS_Batch_Schedule__c`	Stores schedule configuration per definition

Object Relationships

Section titled “Object Relationships”

```
DQS_Definition__c
  ├── DQS_Definition_Detail__c (1:N)
```

```
├── DQS_Batch_Schedule__c (1:1)
└── DQS_Dimension_Result__c (1:N per scan)
    └── DQS_Field_Result__c (1:N, master-detail)
        └── DQS_Metric_Result__c (1:N, master-detail)
```

Custom Metadata Types

Section titled “Custom Metadata Types”

CMTs drive the configuration of capabilities and their evaluation logic. They are package-controlled and not editable by subscribers (except where noted).

CMT	Purpose	Records
`DQS_Capability__mdt`	Defines available quality dimensions	7 (one per capability)
`DQS_Dimension__mdt`	Dimension display configuration	7
`DQS_Metric__mdt`	Metric definitions per capability	Multiple per dimension
`DQS_Input_Configuration__mdt`	Input configurator settings	Per-capability config fields
`DQS_Configuration__mdt`	General app settings	Retention, feature flags

Platform Events

Section titled “Platform Events”

Event	Purpose
`Calculation_Complete__e`	Fired when a scan finishes processing
`DQS_Processing_Error__e`	Fired when an error occurs during batch processing

Feature Parameters

Section titled “Feature Parameters”

Parameter	Purpose
`DQS_AppEnabled`	Controls activation gate — whether the app is licensed and active

Limits & Considerations

Scan Limits

Section titled “Scan Limits”

Each Data Quality Sense activation comes with a fixed number of scans (20 by default). Every scan execution — manual or scheduled — counts toward this quota.

Tracking Your Usage

Section titled “Tracking Your Usage”

The definition card in Insight Studio displays a progress bar showing how many scans have been used. When the limit is approaching, the bar turns red and shows a **Scan limit reached** warning.

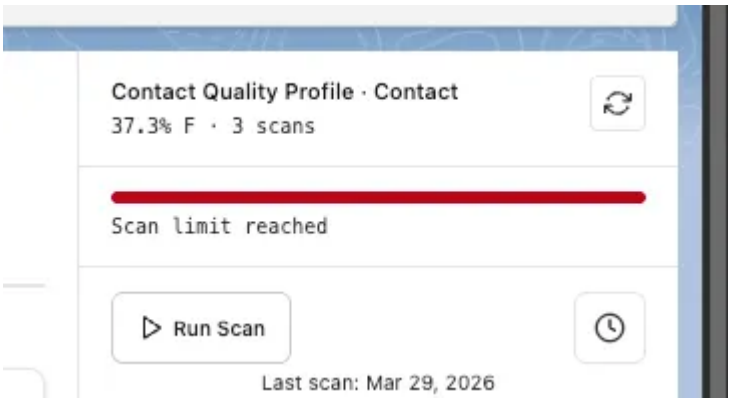

What Happens When the Limit Is Reached

Section titled “What Happens When the Limit Is Reached”

Once all scans are used, attempting to run a new scan triggers an error dialog informing you that the scan limit has been reached.

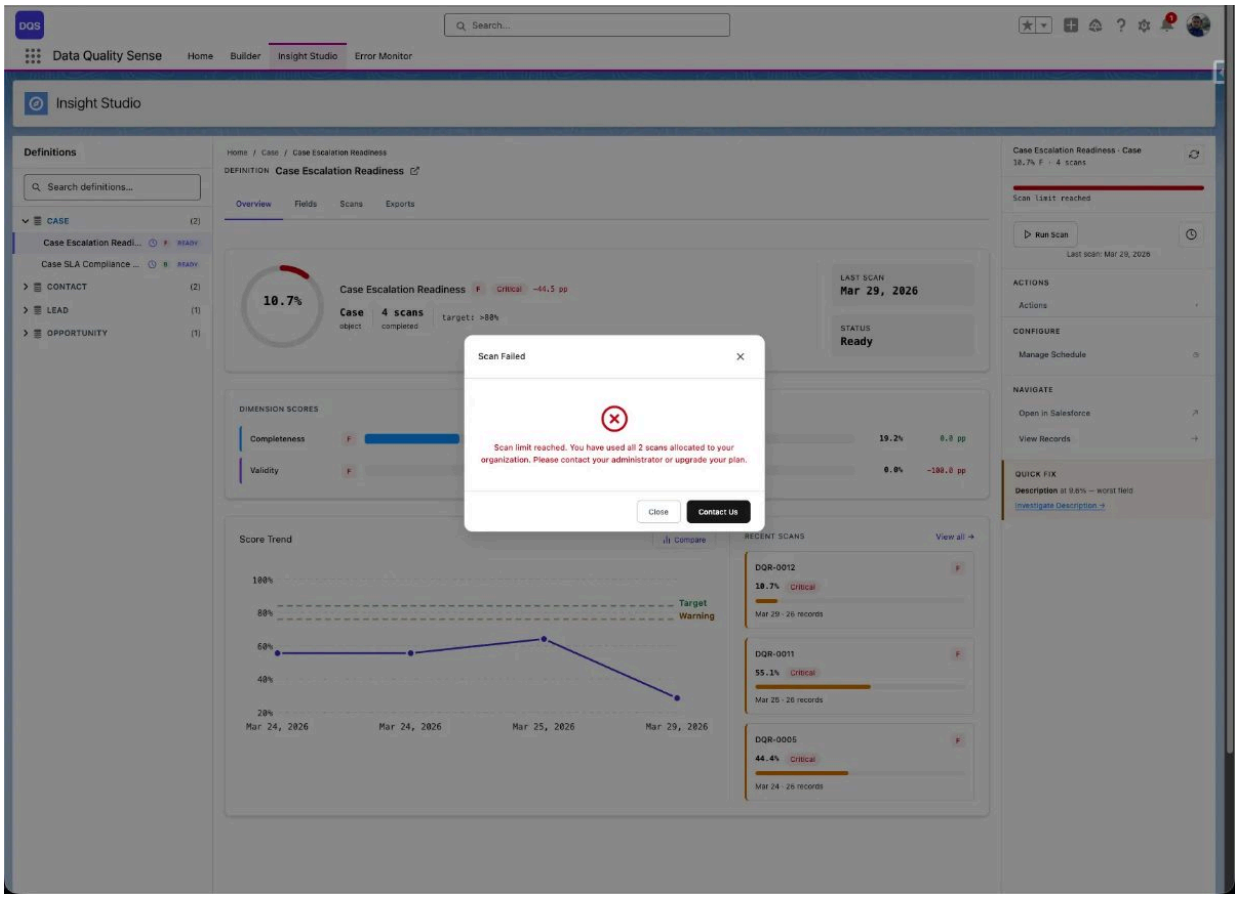

Requesting More Scans

Section titled “Requesting More Scans”

Click the **Upgrade** button in the error dialog to open a pre-filled email requesting additional scans. The email is addressed to hello@tucario.com and includes your Org ID automatically.

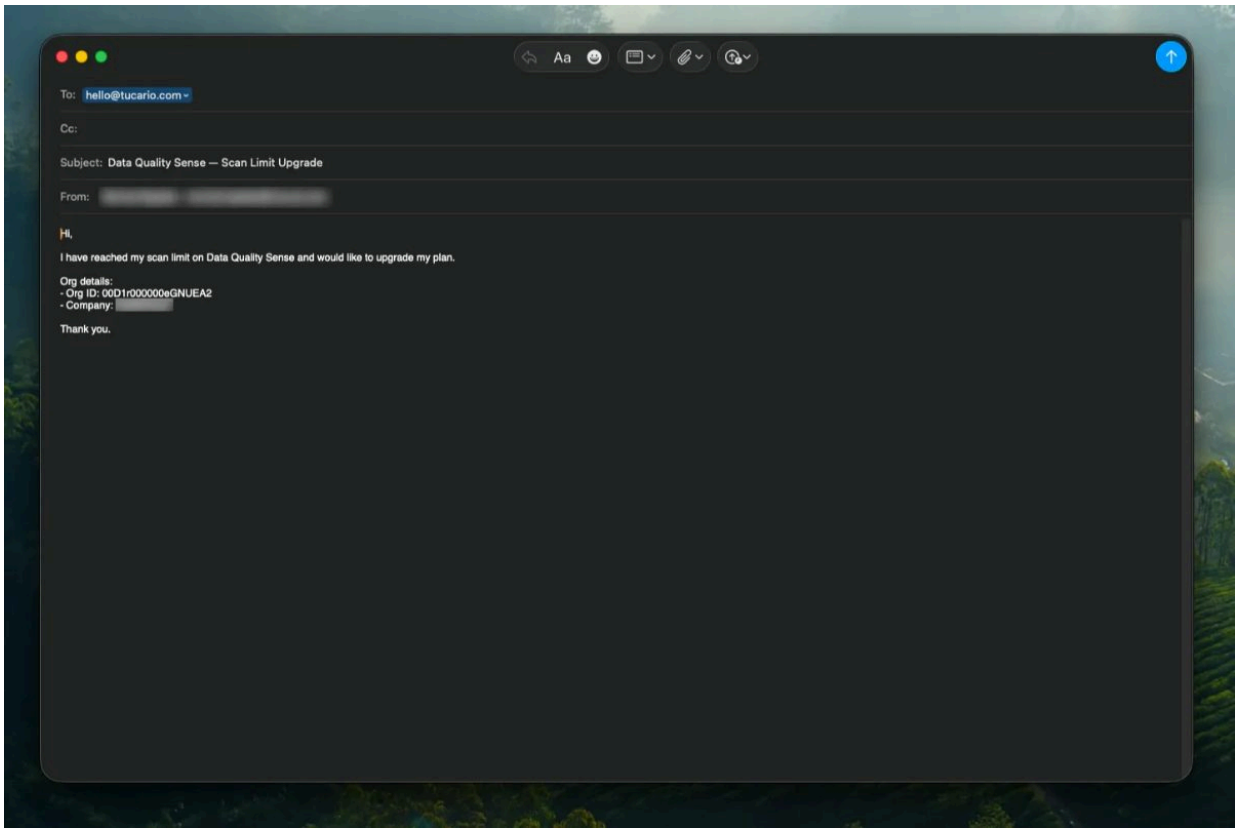

Tip

Scheduled scans also count toward your limit. If you are running low on scans, consider pausing schedules for less critical definitions.

Salesforce Governor Limits

Section titled “Salesforce Governor Limits”

DQS runs entirely within Salesforce and is subject to standard governor limits. The processing engine is designed to work within these constraints.

Batch Processing Limits

Section titled “Batch Processing Limits”

Limit	Salesforce Maximum	DQS Impact
Batch size	2,000 records per chunk	Configurable via DQS settings
Concurrent batches	5 per org	DQS uses 1 batch per scan
SOQL queries per transaction	100	Dynamic queries used efficiently

Limit	Salesforce Maximum	DQS Impact
DML operations per transaction	150	Results batched for efficient writes
Heap size	12 MB (async)	Large text fields may contribute

Scheduling Limits

Section titled “Scheduling Limits”

Limit	Salesforce Maximum
Scheduled Apex jobs	100 per org
CRON triggers	100 per org

⚠Caution

Each DQS scan schedule consumes one scheduled Apex slot. Plan accordingly if your org has many other scheduled jobs.

Performance Considerations

Section titled “Performance Considerations”

Object Size

Section titled “Object Size”

Object Size	Expected Scan Time	Notes
< 10,000 records	Minutes	Fast processing
10,000 – 100,000	10–30 minutes	Normal batch processing
100,000 – 1,000,000	30–60 minutes	Consider off-peak scheduling
> 1,000,000	1+ hours	Schedule during maintenance windows

Number of Fields

Section titled “Number of Fields”

More fields in a definition means more processing per record. A definition with 50+ fields will take longer than one with 10 fields.

Number of Capabilities

Section titled “Number of Capabilities”

Each enabled capability adds a dimension strategy execution per chunk. Enabling all 7 capabilities takes approximately 7x longer than enabling just one.

Best Practices

Section titled “Best Practices”

- **Schedule during off-peak hours** — Minimize impact on users
- **Start small** — Begin with one capability and a few key fields, then expand
- **Monitor batch jobs** — Use Setup → Apex Jobs to check scan progress
- **Use retention policies** — Prevent unbounded result growth
- **Stagger schedules** — Don’t schedule all definitions at the same time

Storage

Section titled “Storage”

Scan results consume Salesforce data storage. Each scan creates:

- 1 Dimension Result per enabled capability
- 1 Field Result per field per capability
- Multiple Metric Results per field

With data retention configured, old results are automatically purged.

Known Issues

Note

These are documented platform constraints or accepted trade-offs, not bugs. We update this list as issues are resolved or new ones are identified.

Actions on Impacted Records

Section titled “Actions on Impacted Records”

Tasks may be created for records whose owner is inactive

Section titled “Tasks may be created for records whose owner is inactive”

When “Create Tasks” is run and a record’s owner is an inactive user, the task for that record fails. Other tasks in the same batch are preserved — only the affected record is skipped. The completion summary reports it as an error.

Workaround: Reassign ownership of records with inactive owners before running Create Tasks, or use the “Assign To” field in the Task modal to assign all tasks to a specific active user.

Chatter posting requires Feed Tracking on the target object

Section titled “Chatter posting requires Feed Tracking on the target object”

The “Post Chatter” action requires Chatter to be enabled on the org and Feed Tracking to be enabled on the specific object being scanned. If Chatter is not enabled, the action returns a clear error message. If Feed Tracking is disabled for the object, the posts fail per-record and the completion summary reports errors.

Workaround: Enable Feed Tracking for the target object in **Setup > Chatter > Feed Tracking** before running Post Chatter.

Mentioning record owners in Chatter posts is slower for large datasets

Section titled “Mentioning record owners in Chatter posts is slower for large datasets”

When “Mention Record Owner” is checked, each Chatter post is created individually rather than in bulk. For large violation sets (500+ records), processing time may be noticeably longer.

Workaround: For large datasets, consider unchecking “Mention Record Owner” for faster processing, then manually notifying owners if needed.

Actions evaluate current data, not the last scan snapshot

Section titled “Actions evaluate current data, not the last scan snapshot”

All actions (Export, Create Tasks, Post Chatter) re-evaluate violations against the current state of the data — not the scan results displayed in the UI. If records have been updated since the last scan, the action may find different violations than what is shown on screen.

Workaround: Run a fresh scan before taking actions if data freshness is critical.

A record may receive multiple tasks or posts across field/dimension combinations

Section titled “A record may receive multiple tasks or posts across field/dimension combinations”

When an action is run across multiple fields and dimensions, a record that violates in multiple combinations will receive one task or Chatter post per combination. Duplicate prevention only operates within a single field+dimension scope.

Example: A record with violations on both Email (Validity) and Phone (Completeness) will receive two separate tasks with different subjects.

Concurrent users may create duplicate actions

Section titled “Concurrent users may create duplicate actions”

If two users run the same action on the same definition at the same time, both deduplication checks pass and duplicate tasks or Chatter posts may be created.

Workaround: Coordinate with team members to avoid running the same action on the same definition simultaneously.

Violations on formula or rollup fields generate non-actionable tasks

Section titled “Violations on formula or rollup fields generate non-actionable tasks”

Scans evaluate all monitored fields, including formula and rollup summary fields. If violations are found on these read-only fields, tasks or Chatter posts are still created — but the record owner cannot directly fix the value since it is calculated.

Workaround: When configuring the scan definition, consider excluding formula and rollup fields from dimensions where violations are not remediable by the record owner.

Export Panel

Section titled “Export Panel”

Export history is limited to 100 files per definition

Section titled “Export history is limited to 100 files per definition”

Definitions with heavy export usage may not show all exported files in the Export History panel — the oldest files are excluded. The files still exist as Salesforce Files and can be found via the Files tab.

Workaround: Delete old exports to make room for newer ones in the panel.

Bulk download may be blocked by browser settings

Section titled “Bulk download may be blocked by browser settings”

The “Download All” action triggers multiple file downloads in sequence. Some browser security settings or extensions may block or warn about multiple simultaneous downloads.

Workaround: Allow multiple downloads for the Salesforce domain in your browser settings, or download files individually.

FAQ

General

Section titled “General”

What Salesforce editions are supported?

Section titled “What Salesforce editions are supported?”

Data Quality Sense works with **Enterprise**, **Unlimited**, and **Developer** editions. It requires Lightning Experience to be enabled.

Does DQS require any external integrations?

Section titled “Does DQS require any external integrations?”

No. DQS is **100% Salesforce-native**. All processing, storage, and visualization happens within your Salesforce org. No data leaves your org.

What is the `dataqualitysens` namespace?

Section titled “What is the dataqualitysens namespace?”

It’s the **managed package namespace** for Data Quality Sense. All custom objects, Apex classes, and LWC components are prefixed with this namespace to avoid naming conflicts with your existing customizations.

Installation

Section titled “Installation”

Can I install DQS in a sandbox first?

Section titled “Can I install DQS in a sandbox first?”

Yes, and we recommend it. Install in a sandbox to evaluate the product, then install in production when ready.

Will installing DQS affect my existing data?

Section titled “Will installing DQS affect my existing data?”

No. DQS only **reads** your existing data during scans. It never modifies your business data. It creates its own custom objects to store scan configurations and results.

How do I upgrade to a newer version?

Section titled “How do I upgrade to a newer version?”

Use the same installation URL. Salesforce detects the existing package and offers an upgrade path. Your definitions and results are preserved.

Configuration

Section titled “Configuration”

How many definitions can I create?

Section titled “How many definitions can I create?”

There is no hard limit on the number of definitions. However, each definition with an active schedule consumes one Salesforce scheduled Apex slot (maximum 100 per org).

Can I scan custom objects?

Section titled “Can I scan custom objects?”

Yes. DQS can scan **any standard or custom SObject** that you have read access to.

Can I change the target object of a definition?

Section titled “Can I change the target object of a definition?”

No. The target object is set at creation time and cannot be changed. Create a new definition for a different object.

Scanning

Section titled “Scanning”

How long does a scan take?

Section titled “How long does a scan take?”

It depends on the number of records, fields, and capabilities. A typical scan of 10,000 records with 3 capabilities takes a few minutes. See Limits for detailed estimates.

Can I run multiple scans simultaneously?

Section titled “Can I run multiple scans simultaneously?”

Each scan uses one Salesforce batch job. Salesforce allows up to 5 concurrent batch jobs per org, so theoretically yes — but we recommend staggering scans to avoid governor limit issues.

What happens if a scan fails?

Section titled “What happens if a scan fails?”

The error is logged in the Error Management Console via platform events. Partial results from completed dimensions are preserved.

Results

Section titled “Results”

How long are results retained?

Section titled “How long are results retained?”

By default, scan results are retained for **30 days** and error logs for **7 days**. These can be adjusted via retention settings.

Can I export results?

Section titled “Can I export results?”

Yes. Insight Studio supports CSV export of violation details for any dimension.

What does a score of 0 mean?

Section titled “What does a score of 0 mean?”

A score of 0 means there was **no data to measure** (zero denominator) — not necessarily that all data is bad. For example, if all fields are null, Completeness can’t calculate a meaningful fill rate.

Support

Section titled “Support”

Where can I get help?

Section titled “Where can I get help?”

Visit dataqualitysense.com for support options, documentation updates, and contact information.