Kenexis® SIS

(formerly Vertigo™)

SIS Lifecycle Management

Revision 3.1 • July 2026

Table of Contents

1 Introduction

1.1 About Kenexis SIS

Kenexis SIS is the Safety Instrumented System (SIS) lifecycle management module of the Kenexis Integrated Safety Suite (KISS). It is a browser-based, multi-user application that supports every stage of the SIS lifecycle defined by IEC 61511 / ISA 84.00.01 — from conceptual design and SIL verification through requirements specification, functional testing, and in-service performance monitoring.

Within one Kenexis SIS study you can:

- **Perform SIL verification** — model each Instrumented Protective Function (IPF), assign failure-rate data and voting to its instruments, and confirm the achieved risk reduction against the selected Safety Integrity Level (SIL) target.
- **Develop the Safety Requirements Specification (SRS)** — maintain general requirements as a structured document and detailed datasheets for every IPF, sensor, logic solver, and final element, plus cause-and-effect diagrams of the trip logic.
- **Track testing and operation** — record proof tests and their outcomes, watch test due dates, log demand events (activations), and compare actual demand rates with the assumptions made during risk assessment.
- **Optimize test intervals** — let the built-in optimizer recommend the longest proof-test intervals that still meet every risk-reduction target.

1.2 Key Terms

Two abbreviations appear throughout this manual. An Instrumented Protective Function (IPF) is an instrumented safeguard — typically a sensor subsystem, a logic solver, and a final element subsystem acting together. A Safety Instrumented Function (SIF) is an IPF that carries a SIL requirement; SIFs are the subset of IPFs governed by IEC 61511. A fuller glossary is provided in Appendix A.

1.3 About Kenexis

Kenexis is an independent engineering consulting firm specializing in the integrity of instrumented safeguards and industrial networks. Kenexis applies skills in risk analysis, reliability engineering, and process engineering to help establish the design and maintenance specifications of instrumented safeguards such as safety instrumented systems, alarm systems, and fire and gas systems. The Kenexis Integrated Safety Suite is the software platform that supports this work.

2 Getting Started

2.1 Signing In

Kenexis SIS runs inside KISS and requires a KISS account. Point your browser at your organization's KISS address and sign in with your username and password, or with a configured single sign-on provider (Azure AD or Okta). Every Kenexis SIS page requires authentication; if your session expires you are returned to the sign-in screen.

2.2 The KISS Manager — Facilities and Studies

After signing in you land on the KISS Manager. Facilities are listed on the left; selecting a facility loads its studies (and any folders) on the right. Kenexis SIS studies are marked with the type "Vertigo". The toolbar above the study list provides New, Move, Copy, Paste, Rename, Delete, Search, Import, Export, and Share.

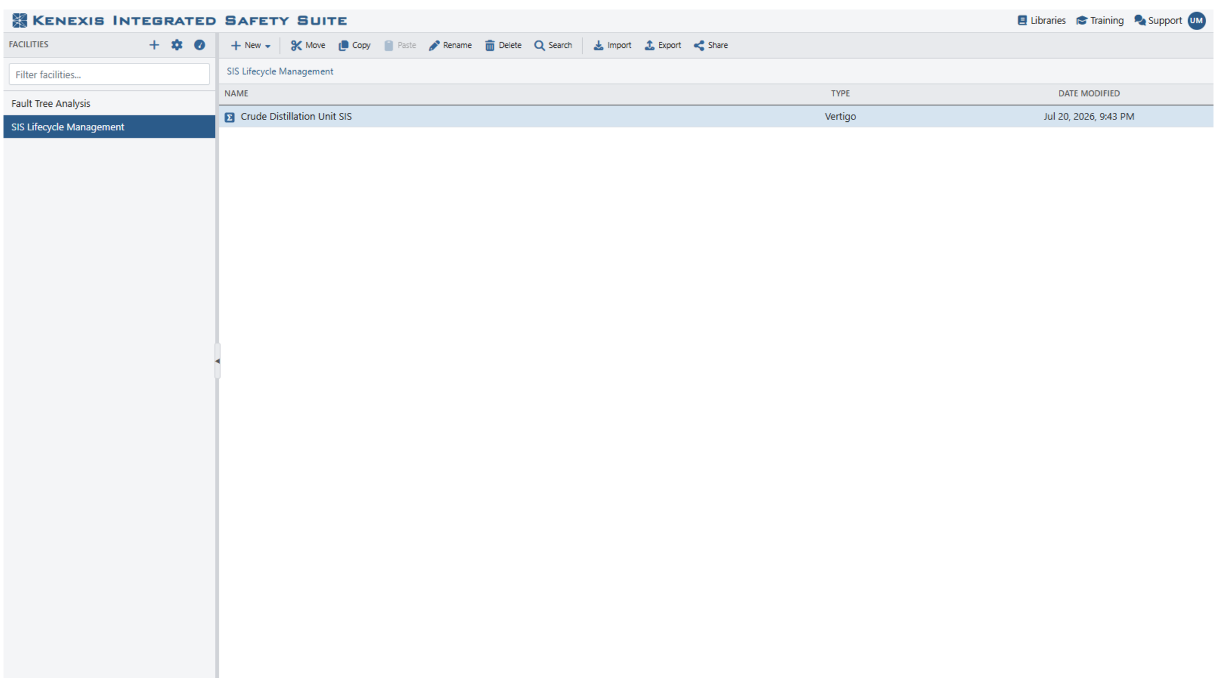

Figure 2-1. The KISS Manager. Selecting a facility lists its studies; the selected study enables the toolbar actions permitted by your role.

2.3 Creating a Kenexis SIS Study

With a facility (or folder within it) selected, choose New › Vertigo and give the study a name. Kenexis SIS creates the study and opens it immediately. A new study starts with the facility's default failure-rate library already linked, so library data is available from the first instrument type you create.

2.4 Managing Studies

- **Copy / Move / Rename / Delete** — duplicate, relocate, or remove studies from the KISS Manager toolbar. Deleting is a soft delete: the study moves to the organization Recycle Bin.
- **Export** — downloads the complete study as a Microsoft Excel workbook (.xlsx), with one worksheet per data category. The workbook can be reviewed or edited in Excel.
- **Import** — uploads a study workbook into the selected facility or folder. KISS detects the study type from the workbook content, assigns the destination facility's default failure-rate library, and runs a full recalculation automatically as part of the import.

Keep the exported column layout

The importer reads worksheet cells by position. Add, edit, and remove rows freely in an exported workbook, but keep the column order exactly as exported.

2.5 The Kenexis SIS Workspace

Opening a study shows the Kenexis SIS workspace. The header carries the study name, a back button to the KISS Manager, and the Study Settings (gear) button. Below the header is the tab bar with the eight workspace tabs:

Table 2-1. The eight Kenexis SIS workspace tabs.

Tab	What it provides
Overview	Study information, documents, recommendations, IPF groups, and the dashboard (Chapter 3).
IPF List	The IPF grid — the main SIL verification surface, with five selectable views (Chapter 4).
Instruments	Sensors, logic solvers, and final elements, with details and testing views (Chapters 5 and 10).
Instrument Types	Failure-rate data records for every device type (Chapter 6).
SRS General Requirements	The structured SRS document editor (Chapter 8).
C&E Diagram	The cause-and-effect matrix of trip logic (Chapter 9).
Reports	Document and spreadsheet outputs (Chapter 13).
Test Interval Optimizer	Automated proof-test-interval recommendations (Chapter 12).

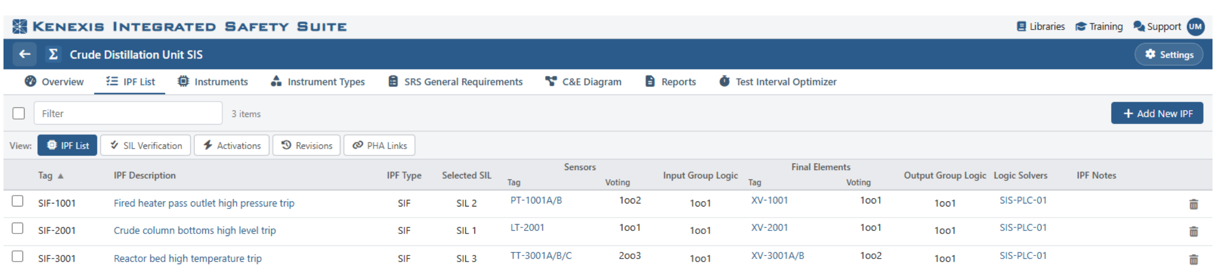

Figure 2-2. The Kenexis SIS workspace, opened on the IPF List tab. Each IPF row shows its sensors, logic solvers, and final elements with their voting.

2.6 Working with Grids

Data grids share a consistent set of controls throughout Kenexis SIS:

- **Filter and count** — the filter box narrows rows by tag, description, type, and similar fields; the live item count sits beside it.
- **Sorting** — click a column header to sort; an arrow indicator shows the current sort column and direction.
- **Selection** — click a row to select it, Ctrl+click to add or remove a row, and Shift+click to select a range. The checkbox column and the header select-all checkbox do the same. Selecting rows reveals the bulk actions available for that grid.
- **Editing** — click the underlined tag or description link in a row to open its details dialog. Dialogs close through their Cancel, Save, or × controls.

2.7 Access Control

Opening a study acquires a license seat and checks your study permissions. With view-only access the workspace shows a Read Only badge and presents all data with the editing controls disabled, so you can review a study while another user retains editing rights.

3 The Overview Tab

The Overview tab collects the study-level records: identification data, the document list, recommendations, IPF groups, and the dashboard. Each lives on its own sub-tab.

3.1 Study Information

The Study Information sub-tab holds the study's identification fields: Study Name, Asset Owner, Project Number, and Status (Planned, Scheduled, In Progress, Draft, Review, Approved, Final, Evergreen, or Archived). Click Update to save; renaming the study here also renames it everywhere it appears in KISS.

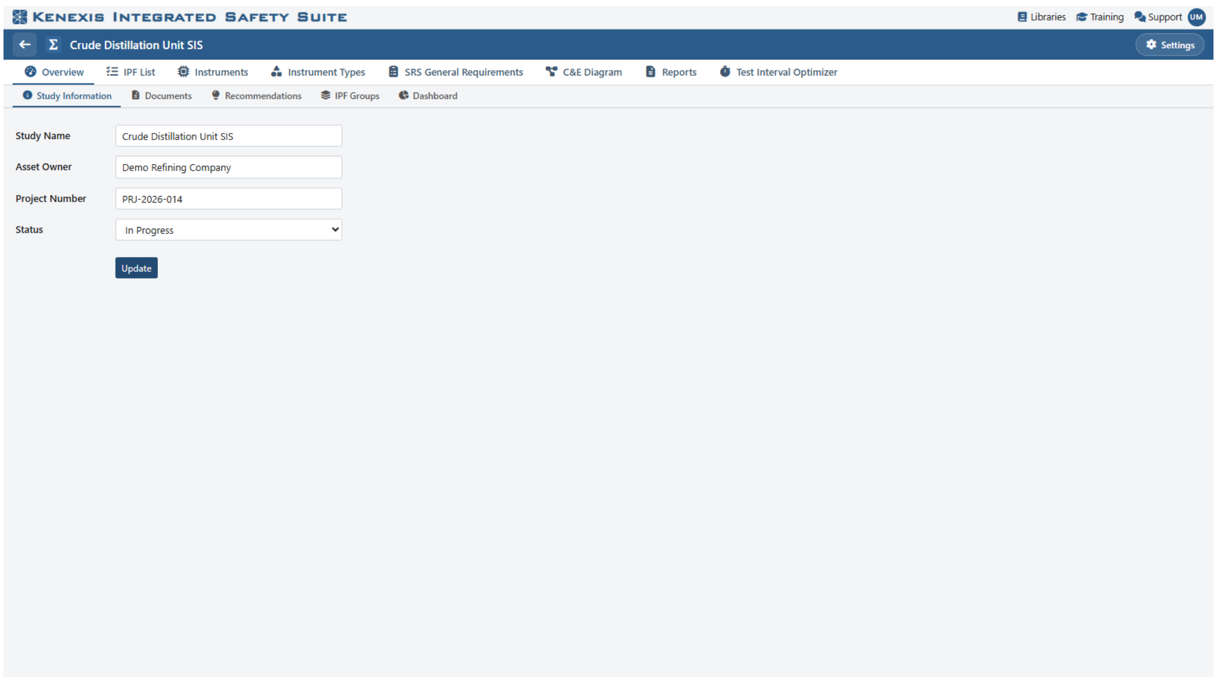

Figure 3-1. The Study Information sub-tab with the study identification fields and the Update button.

3.2 Documents

SIS design work references piping and instrumentation diagrams, cause-and-effect drawings, loop sheets, specifications, safety manuals, and test procedures. The Documents sub-tab tracks these references in one list, so that datasheets elsewhere in the study can point at a document record instead of repeating free text.

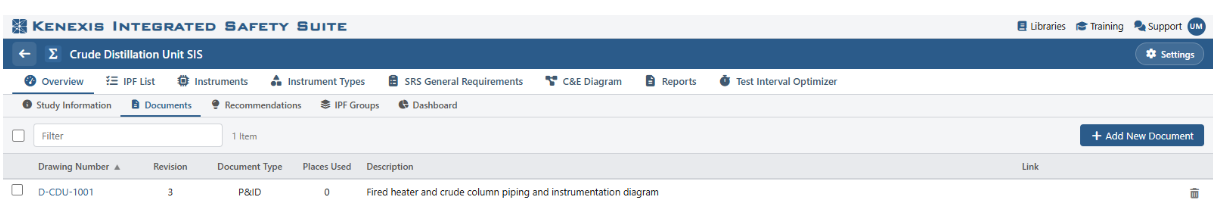

Figure 3-2. The Documents sub-tab. Each record carries a drawing number, revision, document type, description, and hyperlink.

Each document record has a Drawing Number, Revision, Document Type (P&ID, C&E, Loop Sheet, PFD, Safety Manual, Specification, Test Procedure, or Other), Description, and Link. The Link field holds a hyperlink into your document management system, so the source document is one click away without duplicating its content. The Places Used column shows how many instruments reference each document; clicking the count lists them.

3.3 Recommendations

The Recommendations sub-tab lists every recommendation in the study, whether or not it is associated with a specific IPF. Each recommendation carries an automatically assigned number, a Priority (High, Medium, or Low), a Status (from Proposed through Implemented and Closed), a Responsible Party, the recommendation text, and a comment. The Places Used column shows which IPFs include the recommendation and links to each. Recommendations can also be created and linked from inside the IPF dialog (section 4.6).

3.4 IPF Groups

IPF groups collect related functions — by process unit, area, or major equipment item — for filtering and for building cause-and-effect diagrams (Chapter 9). A group is simply a Tag and a Description; assign it to IPFs in the IPF dialog and to instruments on their SRS datasheets.

3.5 Dashboard

The Dashboard sub-tab summarizes the health of the study in six chart cards:

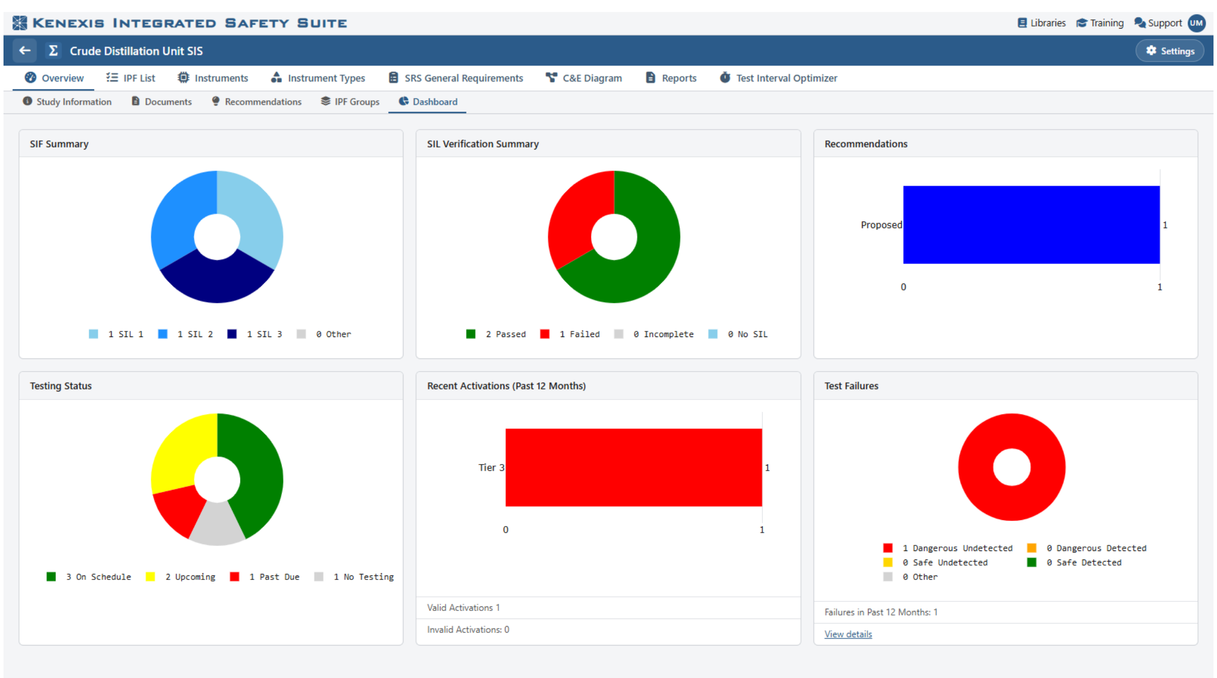

Figure 3-3. The Dashboard. Six cards summarize SIF inventory, verification status, recommendations, testing currency, recent activations, and test failures.

- **SIF Summary** — the inventory of functions by selected SIL.
- **SIL Verification Summary** — how many functions currently pass, fail, are incomplete, or carry no SIL target.
- **Recommendations** — recommendation counts by implementation status.
- **Testing Status** — instruments on schedule, due soon, past due, or without testing data.
- **Recent Activations (Past 12 Months)** — demand events by severity, with valid and invalid activation counts.
- **Test Failures** — functional test failures by failure mode. The View details link opens the Functional Test Failure Summary, which lists each failed test over a selectable window (past year, two years, five years, or all time); clicking a row opens that instrument's testing record.

4 SIL Verification and the IPF List

4.1 Instrumented Protective Functions

SIL verification evaluates the conceptual design of each Safety Instrumented Function: given the selected instruments, their failure rates, voting, and test intervals, does the function achieve its SIL and risk-reduction target? Kenexis SIS performs these calculations continuously — every save of an instrument, instrument type, or IPF recalculates the affected results, so the numbers on screen are always current. The mathematics are documented in full in Chapter 16.

4.2 The IPF Grid and Its Views

The IPF List tab is the primary interface for SIL verification. The View bar switches the grid between five views of the same functions:

Table 4-1. The five views of the IPF grid and the columns each one adds.

View	Columns beyond Tag / Description / Type / Selected SIL
IPF List	Sensors (tag and voting), Input Group Logic, Final Elements (tag and voting), Output Group Logic, Logic Solvers, and IPF Notes. Instrument tags are links that open the instrument's dialog directly.
SIL Verification	IPF Status, Required RRF, Achieved RRF, Minimum Fault Tolerance Satisfied, Max SIL Approved, and the SIL Verification Status light.
Activations	Valid Activations, Operational Time (Years), Activation Rate (per Year), Expected Demand Rate (per Year), and the activation status light (Chapter 11).
Revisions	The latest revision label, date, and the Created By / Checked By / Approved By record.
PHA Links	The count of links to Kenexis Open-PHA safeguards and a per-row Add Link button (section 4.8).

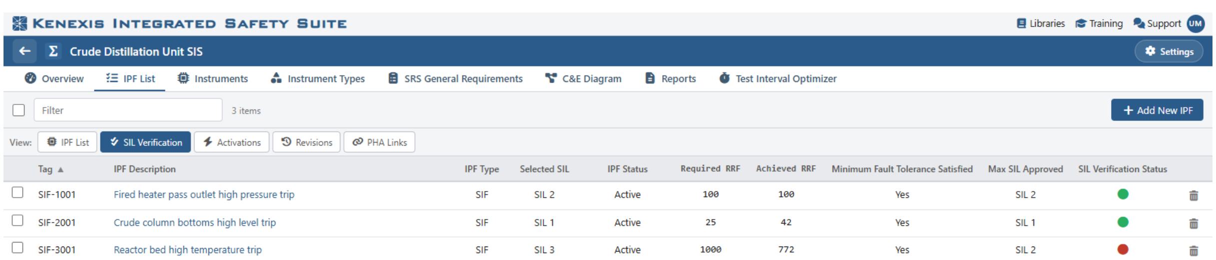

Figure 4-1. The SIL Verification view. The status light is green when the function meets both its risk-reduction target and its fault-tolerance requirement, and red when either is not met.

The SIL Verification Status light reads green when the achieved risk reduction factor meets the required RRF and the minimum fault tolerance is satisfied, and red when either target is missed. A function with no target shows a light only once enough data exists to judge it.

Selecting rows reveals the toolbar actions: Copy (with a choice between creating new instruments or reusing the originals), Copy to Study, Set Status, Convert to Fault Tree (section 4.9), and Delete. The

Revisions view adds bulk Add Revision, Check, and Approve actions that stamp your name and the date on the selected functions.

4.3 The IPF Details Dialog

Open a function by clicking its description (or add one with Add New IPF). The dialog header carries the Tag, the IPF Description, and a live results bar: the achieved-SIL badge, PFD and RRF values, the fault-tolerance verdict, the spurious-trip MTTF, and three gauges showing each subsystem's share of the total PFD.

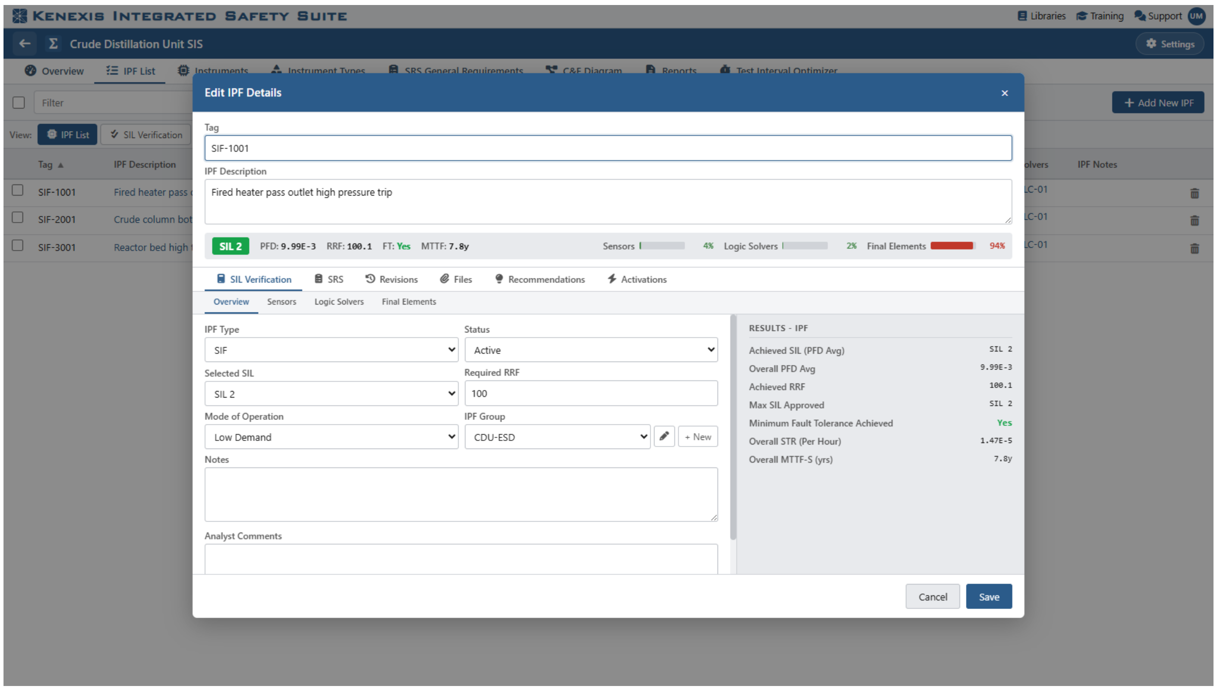

Figure 4-2. The IPF details dialog. The left panel holds the design targets; the Results panel reports the achieved SIL, PFD, RRF, fault tolerance, and spurious-trip performance.

The SIL Verification › Overview panel sets the design targets:

- **IPF Type** — SIF, EPF, MPF, F&G, BPCS, PSV, or No IPF.
- **Status** — Active, Inactive, Future, or Decommissioned.
- **Selected SIL and Required RRF** — choosing a SIL fills the Required RRF with the band floor (10, 100, 1,000, or 10,000); you can then adjust the RRF to the exact target from your risk assessment.
- **Mode of Operation** — Low Demand, High Demand, or Continuous. Continuous mode switches the dialog and the calculations to the per-hour PFH basis of IEC 62061 (section 16.10).
- **IPF Group** — the group assignment, with pencil-edit and + New buttons to manage groups in place.

The Results panel reports Achieved SIL, Overall PFD Avg, Achieved RRF, Max SIL Approved, Minimum Fault Tolerance Achieved, Overall STR, and Overall MTTF-S — the complete verification verdict for the function.

4.4 Assigning Instruments and Group Logic

The Sensors, Logic Solvers, and Final Elements sub-tabs attach instruments to the function. Each sub-tab has a Devices multi-select listing every instrument of that category in the study — tick the ones that serve this function. The + Add button creates a brand-new instrument in place and links it automatically. Closing the device list saves the function and recalculates, so the subsystem results refresh immediately.

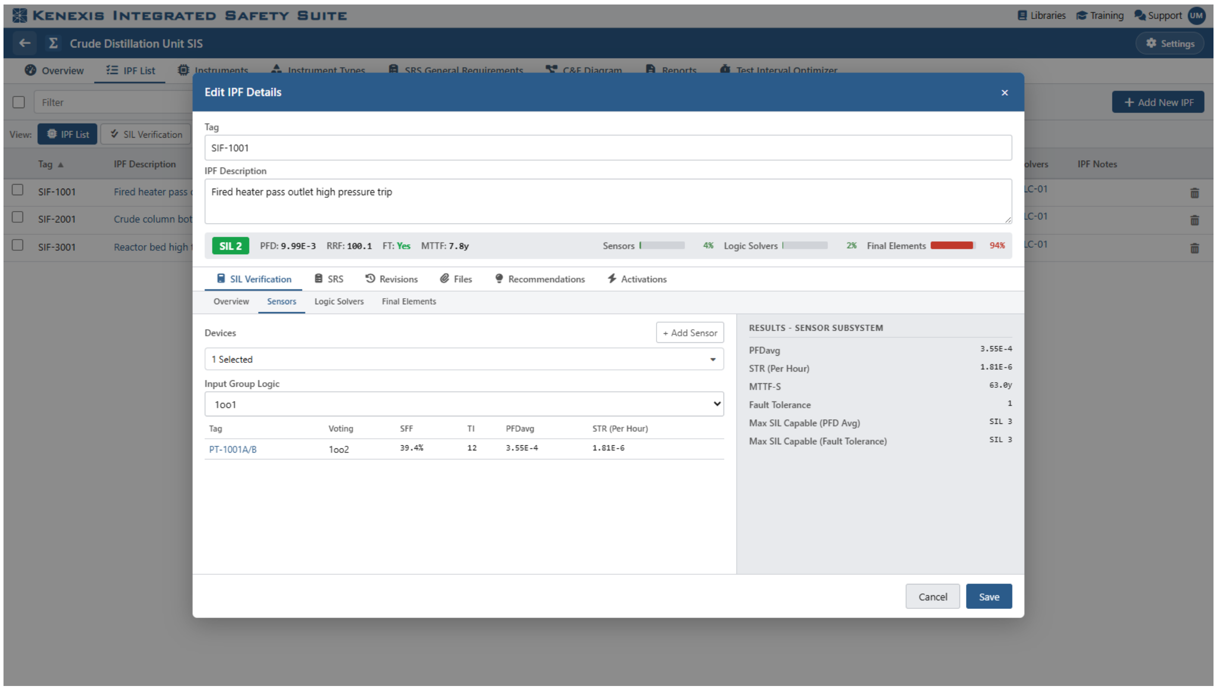

Figure 4-3. A subsystem sub-tab of the IPF dialog. The linked-instrument grid shows each device's voting, SFF, test interval, PFD, and STR; the Results panel rolls them up to the subsystem level.

When a subsystem contains more than one instrument record, the group logic defines how those voted groups combine:

- **1oo1** — a single group makes up the subsystem.
- **1ooX** — any one of several groups can perform the function (redundant groups).
- **XooX** — all groups must operate for the function to succeed.

Sensors use the Input Group Logic setting, final elements the Output Group Logic setting; the logic-solver subsystem combines automatically. Section 16.7 gives the exact formulas.

4.5 Revisions

The Revisions tab of the dialog keeps a per-function change log. Each revision records a label, description, and remarks, stamped with the creator and date; Mark as Checked and Mark as Approved record the reviewer and approver. The same record can be maintained in bulk from the grid's Revisions view.

4.6 Recommendations

The Recommendations tab links study recommendations to the function. Pick existing recommendations from the multi-select, or create one in place with + Add Recommendation — it is numbered automatically and linked to the function as it is created.

4.7 Files

The Files tab attaches supporting files to the function — calculation records, vendor documents, photographs. Chapter 15 describes the file attachment behavior common to IPFs and instruments.

4.8 Links to Kenexis Open-PHA

The PHA Links view of the IPF grid connects functions to the safeguards of a Kenexis Open-PHA study. Add Link opens a picker: choose the facility, the Kenexis Open-PHA study, and the safeguard to link. The link count in the grid opens the list of links, where each entry can be viewed in Kenexis Open-PHA, synchronized, or removed. The Sync dialog compares the safeguard's fields (description, tag, type, SIL, PFD or IPL credits, test interval, response time) with the IPF side by side and copies any field from the PHA record onto the function with one click.

4.9 Converting an IPF to a Fault Tree

With exactly one function selected, Convert to Fault Tree builds a complete Kenexis Fault Tree fault-tree study from the function — its sensors, logic solvers, and final elements become gates and basic events with the same failure data — runs the Kenexis Fault Tree calculation, and offers a View link to open the new study. This provides a fully worked fault-tree model wherever a function needs analysis beyond the simplified equations.

5 Instruments

5.1 The Instrument Grids

The Instruments tab holds one grid per category — Sensors, Logic Solvers, and Final Elements. An instrument record represents a voted group of identical devices: a single transmitter, or a matched pair voting 1oo2, is one record with the voting set accordingly.

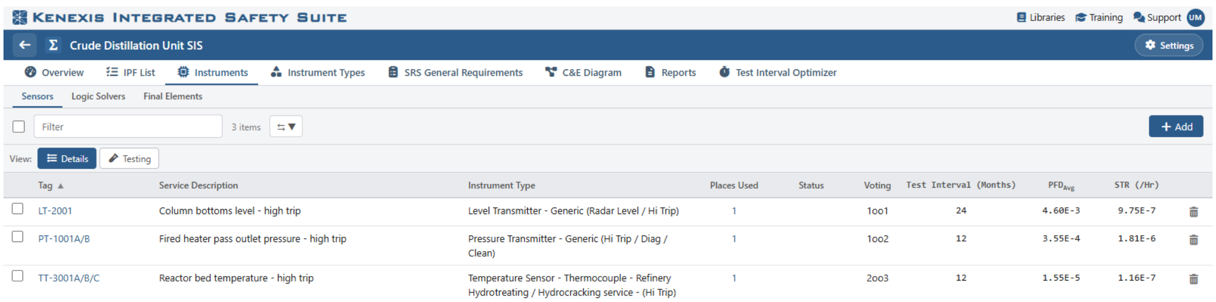

Figure 5-1. The Sensors grid in Details view: tag, service description, instrument type, places used, status, voting, test interval, and the calculated PFD and STR.

The Details view shows each instrument's tag, service description, instrument type, places used (the IPFs that use it — click the count to see them), status, voting, test interval, and the calculated PFD and STR. For sensors and final elements a toggle reveals the interface columns (process connection and interfaces). The Testing view is described in Chapter 10.

Selecting rows reveals Copy, Set Status (Active, Pending, Bypassed, Out of Service, Decommissioned), Update Test Interval (set one interval on many instruments at once), and Delete.

5.2 The Instrument Details Dialog

Click an instrument's tag to open its details dialog. The SIL Verification tab pairs the input panel on the left with a live results panel on the right; every change updates the results immediately, before you save.

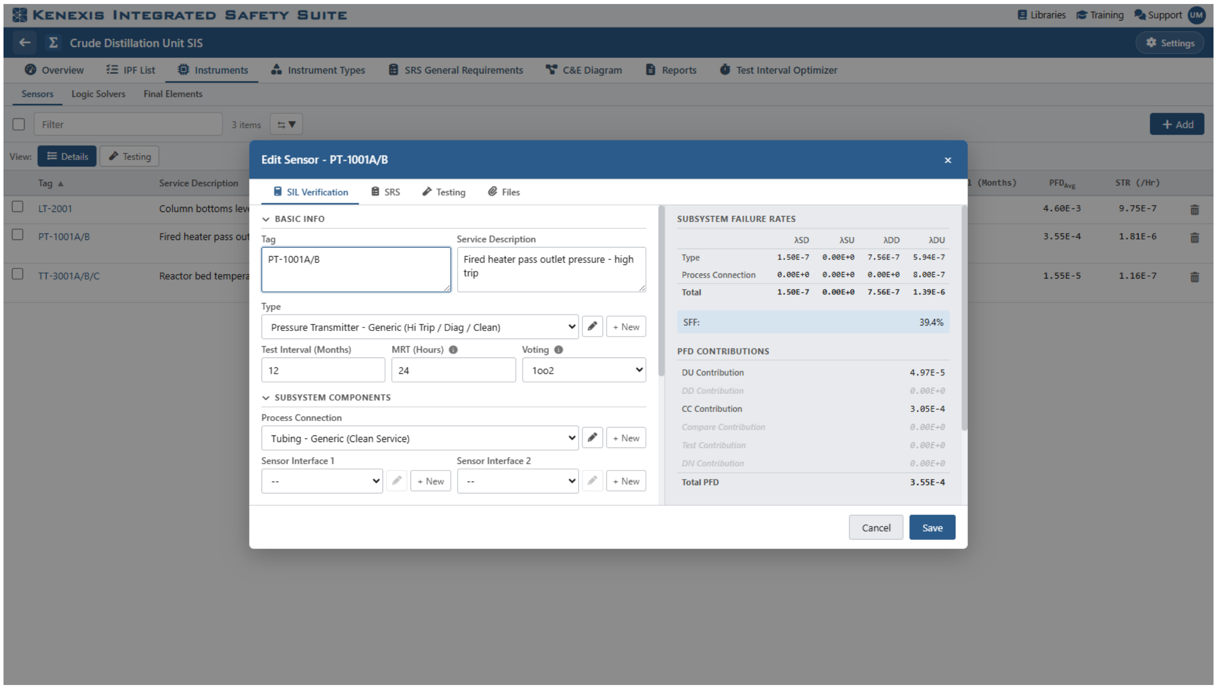

Figure 5-2. The instrument details dialog. Inputs on the left; the failure-rate roll-up, SFF, PFD contributions, and STR contributions update live on the right.

Basic Info

Tag, Service Description, the instrument Type (chosen from the study's instrument types — the pencil edits the type, + New creates one in place), Test Interval in months, MRT (mean repair time) in hours, and Voting.

Voting

Voting expresses how many of the installed devices must vote to trip: 1oo1, 1oo2, 2oo2, 1oo3, 2oo3, 3oo3, 1oo4, 2oo4, or Custom (MooN) with any M of N. M is the number of devices that must act; N is the number installed. The voting drives both the probability of failure on demand and the spurious trip rate (sections 16.3–16.5).

Subsystem Components

A sensor can include a Process Connection and up to two Sensor Interfaces; a final element up to two Final Element Interfaces. Each component contributes its own failure rates, which are summed with the device type's rates in the Subsystem Failure Rates table (section 16.2). Every component dropdown has pencil-edit and + New buttons.

Advanced Parameters

Five options refine the calculation; each enables the inputs it needs:

Table 5-1. Advanced parameters and the inputs each one enables.

Option	Inputs it enables	Effect
Common Cause	Beta Factor (0–1)	Adds a common-cause contribution to PFD and STR for redundant voting. The calculator button opens the Beta Factor Estimator (Chapter 7).
Trip on Detected Failure	Diagnostic Interval (when off)	On: detected failures produce a trip vote and contribute to the spurious trip rate. Off: they are alarmed and repaired, contributing to PFD across the diagnostic interval.
Online Testing	Test Duration (hours)	Models the unavailability of a device that is bypassed while it is proof tested in service.
Imperfect Testing	Proof Test Coverage (0–1), Useful Life (years)	Splits dangerous undetected failures into a tested part and a never-detected part that persists until the device is replaced at the end of its useful life.
External Comparison	Comparison Coverage (0–1), Comparison Interval (hours)	Credits an external comparison (for example, cross-checks against another measurement) that reveals additional dangerous failures.

5.3 The Live Results Panel

The right-hand panel reports, as you type:

- **Subsystem Failure Rates** — λSD, λSU, λDD, and λDU for the type and each component, with the totals and the resulting Safe Failure Fraction (SFF).
- **PFD Contributions** — the DU, DD, CC, Compare, Test, and DN contributions and the Total PFD. Rows tied to an option that is switched off are grayed out.
- **STR Contributions** — the SU, DD, and CC contributions to the spurious trip rate per hour, and the total.

Chapter 16 defines every one of these quantities and the formulas behind them.

6 Failure Rate Data and Instrument Types

6.1 The Instrument Type Concept

Failure-rate data lives on instrument types, not on individual tagged devices. Define one type per make and model — for example a generic high-trip pressure transmitter — and apply it to every instrument of that kind. The data is entered once and used everywhere, and updating the type recalculates every instrument and function that uses it.

Types are organized in six categories on the Instrument Types tab: Process Connections, Sensor Interfaces, Sensor Types, Logic Solver Types, Final Element Interfaces, and Final Element Types.

6.2 The Instrument Types Grid

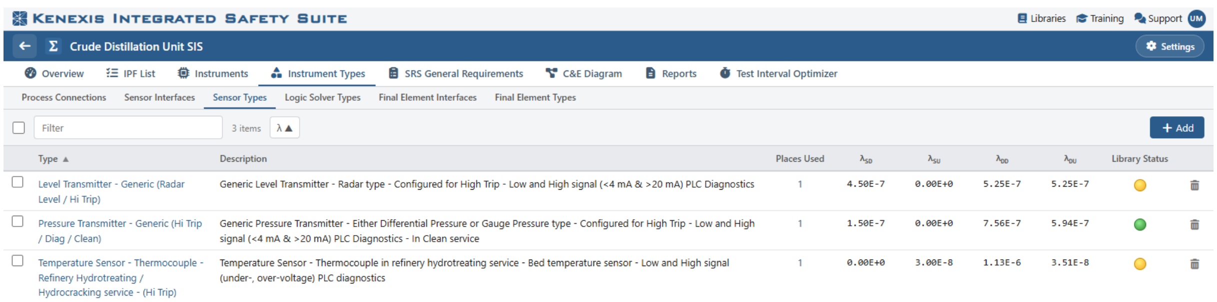

Figure 6-1. A category grid on the Instrument Types tab, showing the four failure-rate splits and each type's library status.

Each row shows the type, description, places used, its failure-rate data, and a Library Status dot. The λ toggle in the toolbar switches the data columns between the four failure-rate splits (λSD, λSU, λDD, λDU) and the summary form (failure rate, percent safe, safe coverage, dangerous coverage). Selected rows can be duplicated, deleted, or — when you have write access to a failure-rate library — added to that library for reuse in other studies. Types that are in use by an instrument are protected from deletion.

6.3 Creating Types from a Library

Click Add to create a type. The Select Instrument Type picker searches the study's failure-rate library as you type; choosing an item copies its complete data — failure rates, architectural constraint type, and data reference — into the form. Choosing Custom Component starts a blank form for data you supply yourself.

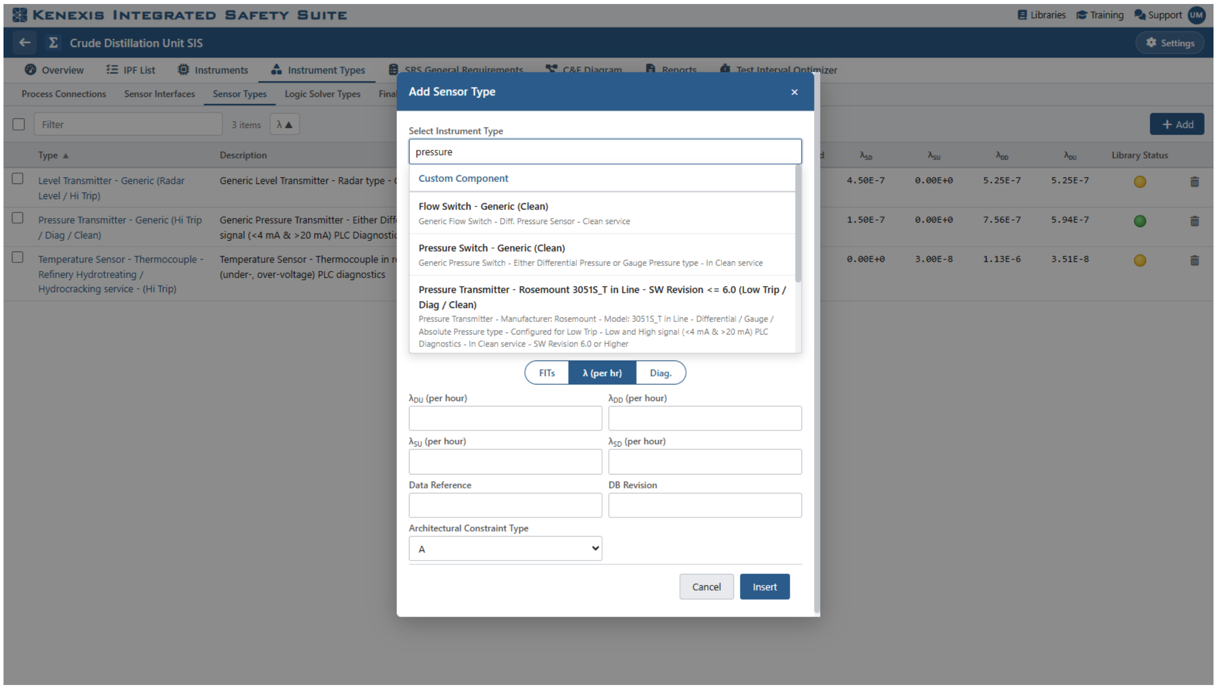

Figure 6-2. Adding a type. The picker filters the connected failure-rate library; selecting an item fills the form with its data.

Every Kenexis SIS study is connected to a failure-rate library — the Kenexis Standard library by default, which Kenexis engineers maintain with current industry data. Organizations can also build their own libraries and select them per study in Study Settings (Chapter 14).

6.4 Failure-Rate Entry Modes

For custom data, the form offers three equivalent entry modes; pick whichever matches your data source. The mode you choose is remembered for next time.

- **FITs** — the four splits in failures per 10^9 hours.
- **λ (per hr)** — the four splits as rates per hour, in scientific notation.
- **Diag.** — a single overall failure rate plus Percent Safe, Safe Coverage, and Dangerous Coverage; Kenexis SIS derives the four splits (section 16.2).

Each type also carries its Architectural Constraint Type (A for simple devices with well-understood failure modes, B for complex microprocessor-based devices), which feeds the fault-tolerance assessment (section 16.9). Logic solver types can instead record a SIL certification level (SIL 2 or SIL 3).

6.5 Library Status and Synchronization

The Library Status dot keeps study data honest against its source library: green means the type matches the current library data; yellow means the library has moved on — click the dot to open a side-by-side comparison of study and library values and, if you choose, synchronize the study to the library; gray means the type is custom or its library item is no longer available.

6.6 Black Box Models and Fault-Tree Links

Some subsystems are richer than the simplified equations can express — complex valve arrangements, cross-connected sensor networks, or anything you have already modeled in detail. A Black Box Model (BBM) type carries its results directly: you provide the PFDavg, the STR, the hardware fault tolerance, and the maximum SIL approved, and instruments using the type adopt those values as-is.

Each BBM value can either be typed in or linked to a Kenexis Fault Tree fault-tree study with Link to Fault Tree. A linked value is read-only and shows the source study; every time that fault tree recalculates, Kenexis Fault Tree pushes its top result into the type and Kenexis SIS recalculates everything that depends on it. The linked value therefore always reflects the fault tree's latest calculation. Combined with IPF-to-fault-tree conversion (section 4.9), this closes the loop between the two modules in both directions.

7 The Beta Factor Estimator

Redundant devices share an environment, a design, and a maintenance program, so a fraction of their failures strike all channels at once. The beta factor is that fraction, and it usually dominates the PFD of a redundant subsystem. The Beta Factor Estimator turns the qualitative checklist method of IEC 61508-6 Annex D into a guided scoring dialog, launched from the calculator button beside the Beta Factor field of any instrument.

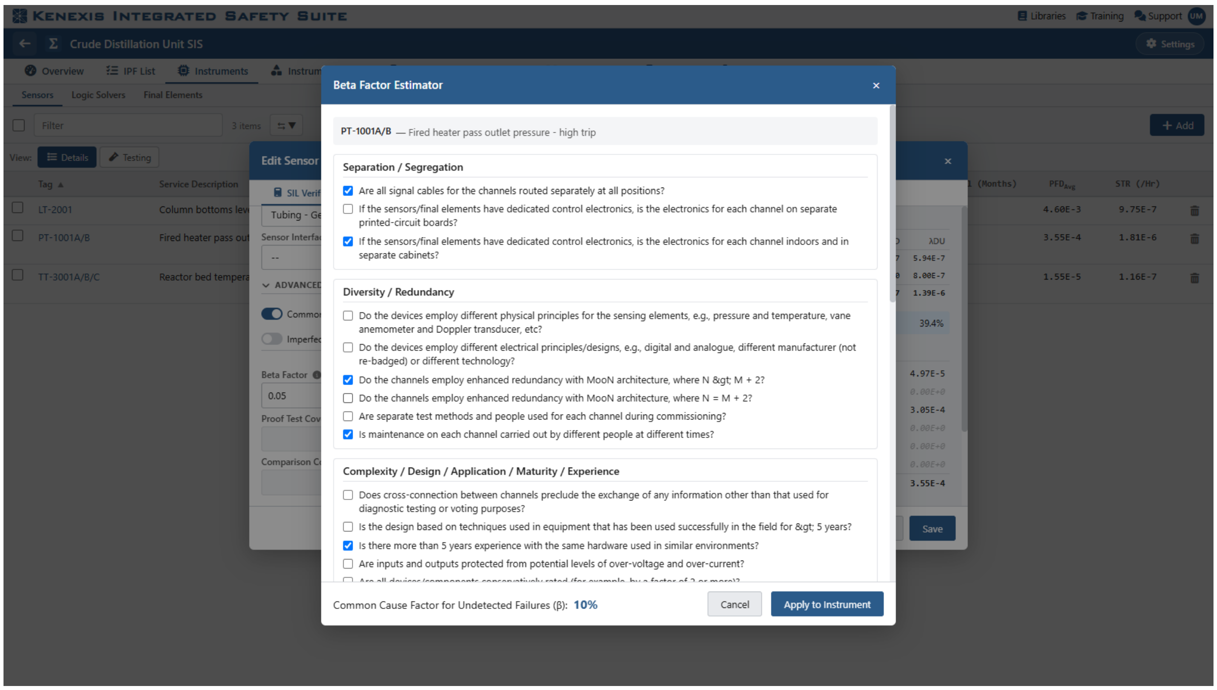

Figure 7-1. The Beta Factor Estimator. Checking the defenses that apply to the installation scores the design; the resulting common-cause factor updates live at the bottom.

The checklist covers eight categories of defense against common cause: separation and segregation, diversity and redundancy, design complexity and maturity, assessment and analysis, procedures and human interface, competence and safety culture, environmental control, and environmental testing. Field devices answer one question set; logic solvers can switch to a longer set tailored to programmable systems. Each checked defense adds points, and the total score maps to a beta factor (section 16.13 gives the exact bands).

Apply to Instrument writes the estimated beta onto the instrument, records the answers so the assessment can be revisited later, and switches Common Cause on.

8 Safety Requirements Specification

8.1 The Two Halves of the SRS

IEC 61511 requires a Safety Requirements Specification documenting what every safety instrumented function must do. Kenexis SIS splits the SRS the way most operating companies do:

- **General requirements** — the rules that apply across the whole SIS, written once as a structured document (section 8.4).
- **Datasheets** — the per-item specifics, captured as fields on each IPF, sensor, logic solver, and final element (sections 8.2 and 8.3).

8.2 Tracked Fields

The SRS datasheet field set is deliberately a superset — far more fields than any one organization uses. Study Settings › SRS Tracked Fields (Chapter 14) chooses which fields appear on the forms and reports for this study, so datasheets show exactly the fields your specification tracks.

8.3 Datasheets

The IPF dialog's SRS tab collects the function-level requirements: basic data, HAZOP and LOPA references, the logic and operation description, process safety time and response times, testing requirements, and notes.

Each instrument's SRS tab carries the device-level datasheet. All categories record status, the device selection basis (Prior Use, Proven In Use, or IEC 61508 Compliant), and references to the P&ID, C&E, and loop sheet drawings — each drawing field links to the study document list, with + New creating a document record in place. Category-specific sections then follow:

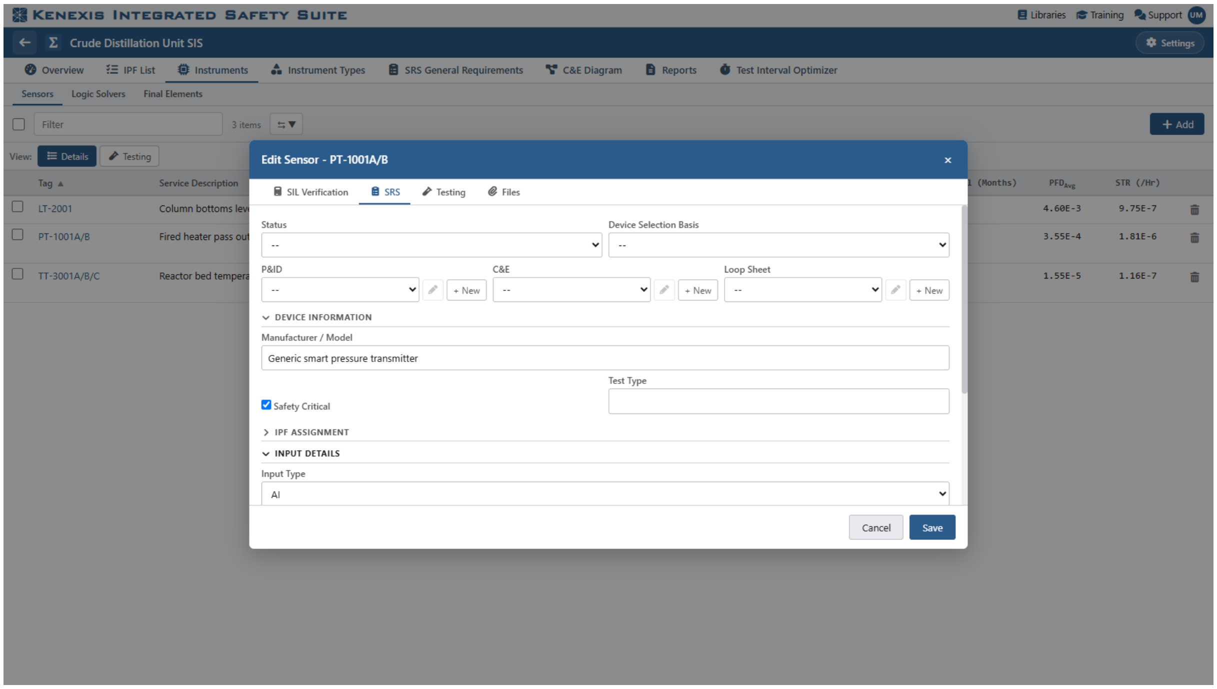

Figure 8-1. A sensor SRS datasheet. Sections expand to show the fields selected in the study's tracked-field settings; drawing references link to the document list.

- **Sensors** — device information, IPF group assignment, input details (range, engineering units, trip setpoint and tolerance, response time, fault response), deviation alarms, fault and shutdown alarms, bypasses, and notes.
- **Logic solvers** — device information, the primary HMI, power supplies, and fault alarms.
- **Final elements** — device information, IPF group assignment, the SIF action and success criteria, output details, tight shutoff requirements, position feedback, fault alarms, reset and manual trip provisions, bypasses, and notes.

Bypass provisions are recorded per device: a maintenance bypass and an operational bypass, each with its location, bypass tag, alarm tag, and service description — so authorized bypass points are part of the specification itself.

8.4 The General Requirements Editor

The SRS General Requirements tab is a structured document editor. The outline pane manages numbered sections; each section holds numbered clauses and sub-clauses with rich text — bold, italics, lists, tables, and images. Requirement numbers renumber automatically as you reorganize, edits save automatically a moment after you stop typing, and Import from Library loads a requirements template from a KISS library as a starting point.

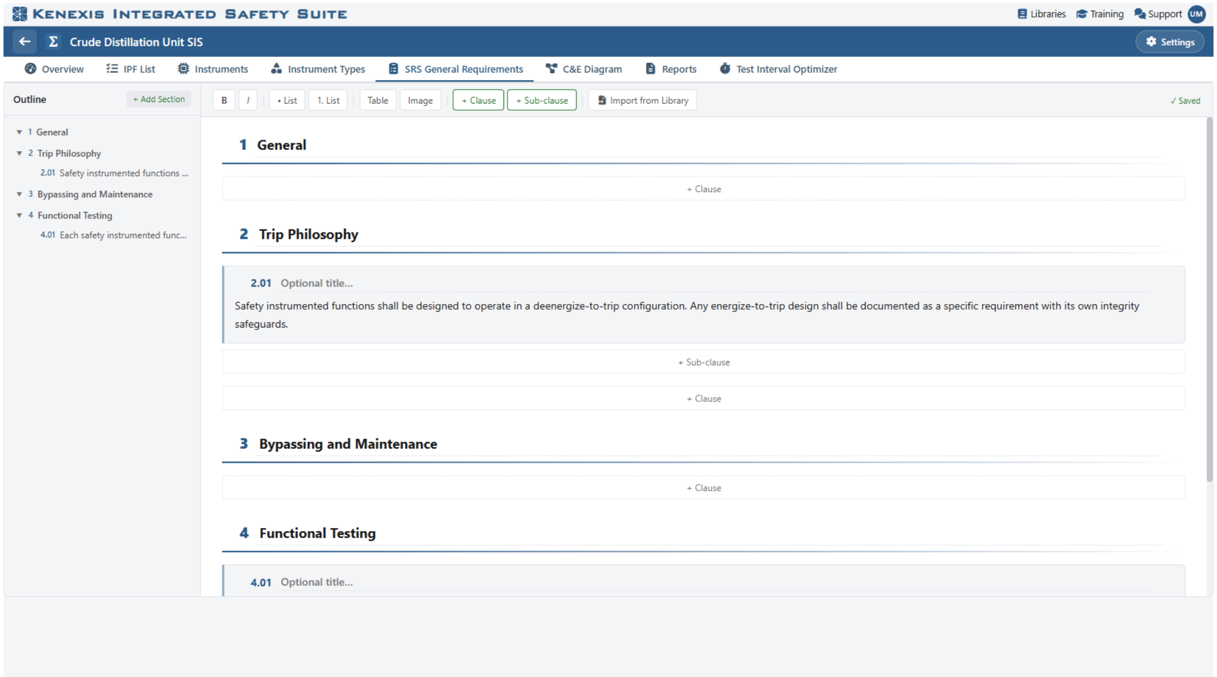

Figure 8-2. The SRS General Requirements editor: the outline on the left, the numbered document with its formatting toolbar on the right.

9 Cause and Effect Diagrams

The C&E Diagram tab presents the trip logic as a live matrix: sensors as rows (causes), final elements as columns (effects), and the intersections describing what each sensor activation does to each final element. The matrix is built automatically from the study's instruments, filtered by IPF group.

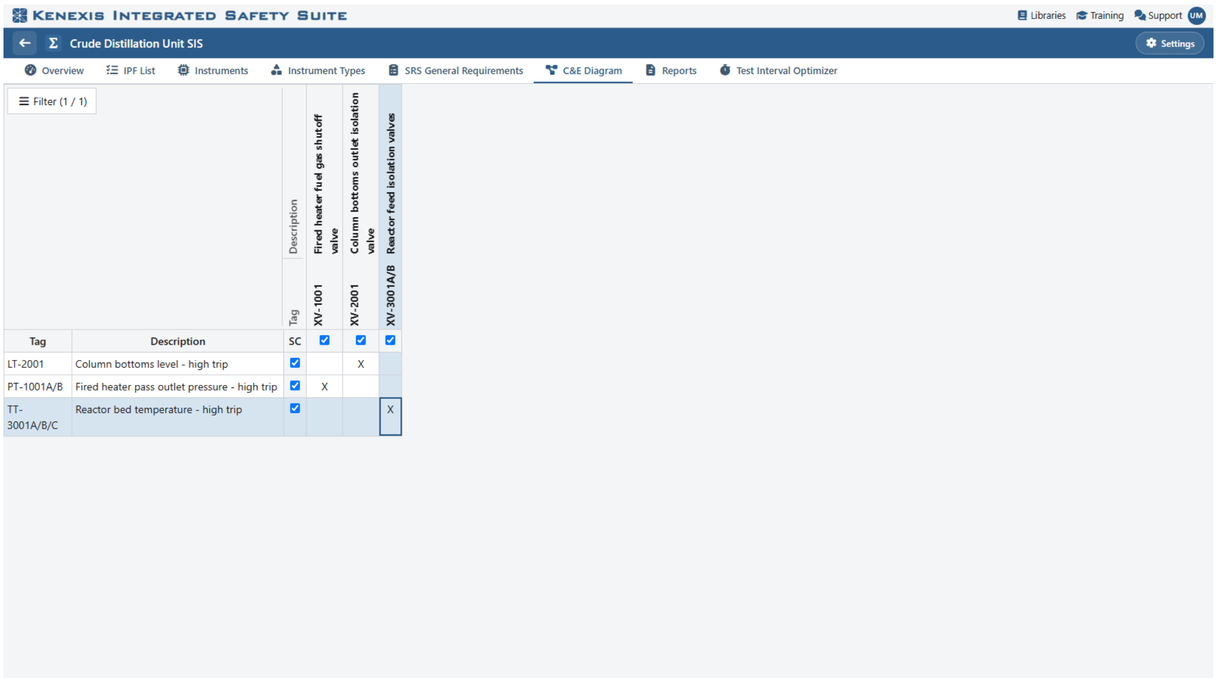

Figure 9-1. The cause-and-effect matrix. Sensor rows and final-element columns come from the selected IPF groups; intersection marks define the trip actions.

The filter panel selects which IPF groups appear — instruments join the matrix through the IPF group on their SRS datasheet. The SC checkboxes toggle each instrument's Safety Critical flag and save immediately, and the tag links open the instrument dialogs.

Intersections edit like a spreadsheet: click a cell to select it, type to start editing, Enter to commit, Escape to cancel, and arrow keys to move between cells. An entry holds up to five characters — commonly an X, an action code such as OPEN or STOP, or a note reference that points into the general requirements. Clearing a cell removes the intersection. The Cause & Effect Diagrams report (Chapter 13) exports the same matrix to Excel, one worksheet per IPF group.

10 Functional Test Tracking

10.1 The Testing View

Proof testing keeps the PFD calculation honest — the calculated risk reduction assumes every dangerous undetected failure is found at the test interval. The Testing view of each instrument grid turns those intervals into a live schedule.

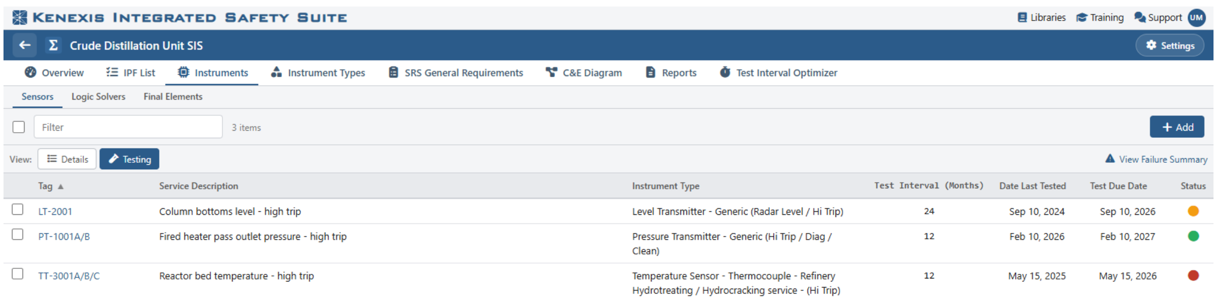

Figure 10-1. The Testing view of an instrument grid: test interval, date last tested, test due date, and the status light for each instrument.

For each instrument the view shows the Test Interval, Date Last Tested (the most recent recorded test, or the commission date before any tests exist), the Test Due Date (last test plus interval), and a status light:

Table 10-1. Test due-date status lights.

Light	Meaning
Green	The next test is due more than six months out.
Yellow	The next test is due within the next six months (182 days).
Red	The test due date has passed.
None	The instrument is decommissioned (shown as O.O.S.), or there is not yet enough data to compute a due date.

10.2 Recording Tests

Open an instrument and choose its Testing tab to maintain the record. The header holds the Commission Date and Decommission Date; below it, Add Test records each proof test: the date, whether it passed, the work order number, and notes. For a failed test you also record which component failed (the device type or one of its interfaces) and the failure mode — Dangerous Detected, Dangerous Undetected, Safe Detected, or Safe Undetected.

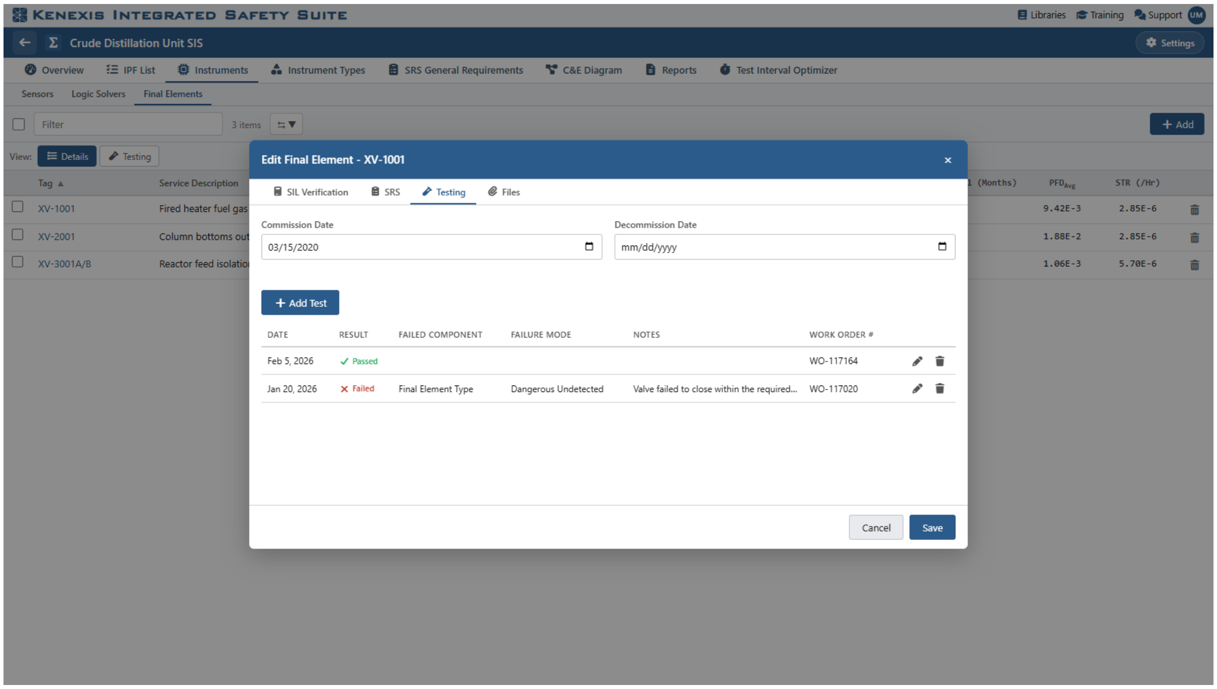

Figure 10-2. An instrument's test history. Failed tests carry the failed component and failure mode, feeding the failure-rate-from-testing reports.

Recorded failures power two follow-ups: the Test Failures dashboard card (section 3.5) and the Failure Rates Based on Testing reports (Chapter 13), which compute actual in-service failure rates you can compare against the library data used in the SIL verification.

10.3 Bulk Test Entry

A test campaign rarely touches one device at a time. In the Testing view, select up to fifty instruments and choose Add New Test to record the same test event — date, result, work order, and failure details if any — against every selected instrument at once.

11 Activation Tracking

11.1 Why Track Activations

Every real demand on a safety function is evidence about the demand-rate assumption made when its SIL was selected. Kenexis SIS records these events per function, computes the actual demand rate, and flags functions whose reality has outrun their risk assessment; API RP 754 severity tiers on each event support process-safety performance reporting.

11.2 The Activations View

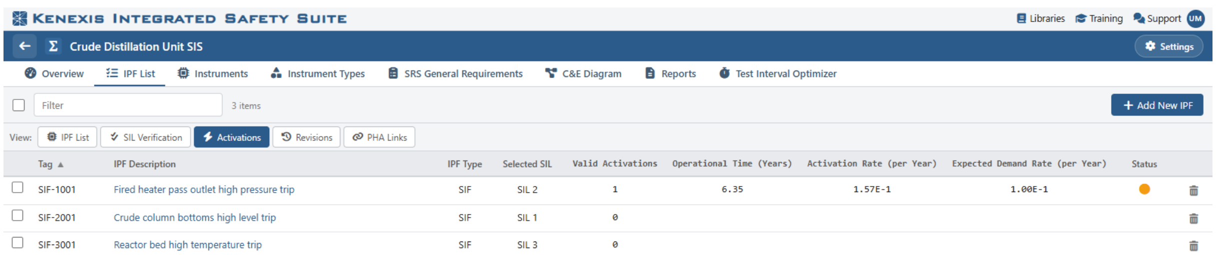

Figure 11-1. The Activations view of the IPF grid: valid activations, operational time, the computed activation rate, the expected demand rate, and the comparison light.

For each function the view shows the count of valid activations, the operational time since commissioning, the resulting Activation Rate per year, and the Expected Demand Rate from the risk assessment. The light — shown once a commission date and expected rate are recorded — is green while the actual rate stays below the expected rate, and yellow once it reaches it: the signal to revisit the risk assessment or the SIL target.

11.3 Recording Events

The Activations tab of the IPF dialog holds the record. The header sets the Commission Date, Decommission Date, and Expected Demand Rate per year. Add New Event records each activation: the date and time, whether it was a valid demand (a genuine process demand rather than a spurious activation), the API RP 754 severity tier, and notes; an invalid event also records the reason. Historian-collected events are marked as such, and only valid events count toward the activation rate.

12 The Test Interval Optimizer

Testing costs money; testing too rarely costs risk reduction. The Test Interval Optimizer searches for the longest proof-test intervals at which every selected function still meets its required risk reduction factor.

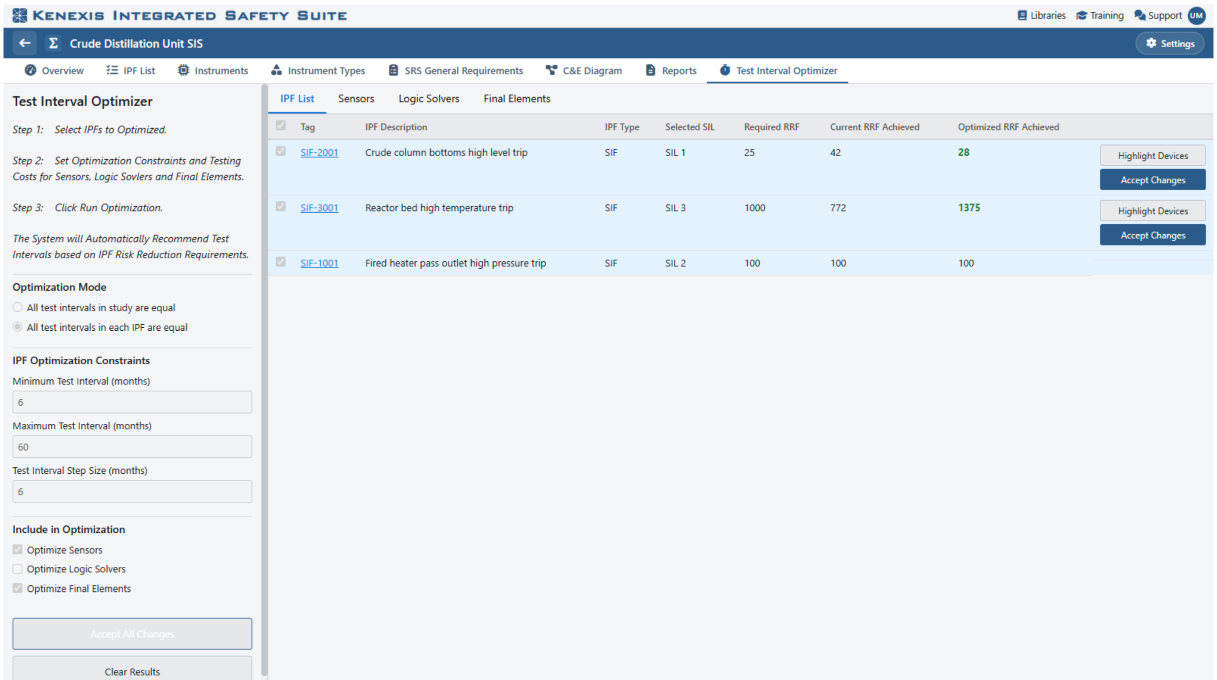

Figure 12-1. The Test Interval Optimizer. Constraints on the left; the results grids compare current and optimized values, with green marking targets met and red marking targets missed.

Set up a run in three steps:

1. Select the functions to optimize in the IPF List grid (functions need a Required RRF to participate).
2. Choose the optimization mode and constraints: the minimum and maximum test interval and the step size, in months, plus which device categories (sensors, logic solvers, final elements) the optimizer may adjust.
3. Click Run Optimization.

The two modes differ in how intervals are shared:

- **All test intervals in study are equal** — one common interval for every optimized device, stepped down from the maximum until every selected function passes.
- **All test intervals in each IPF are equal** — each function's devices share an interval, optimized per function; devices shared between functions keep the interval required by the most demanding one.

Results appear side by side: each function's current and optimized achieved RRF, and each device's current and optimized interval, with the functions it serves rolled up beneath it. Highlight Devices jumps from a function to its affected devices. Nothing changes until you accept — per device, per function, or Accept All Changes — which writes the intervals to the instruments and recalculates the study. Skipped items are listed with the reason.

13 Reports

The Reports tab generates the study's documents. Pick a report from the category tree on the left, tick the items to include in the selection grid, and click Generate Report; the file downloads when it is ready.

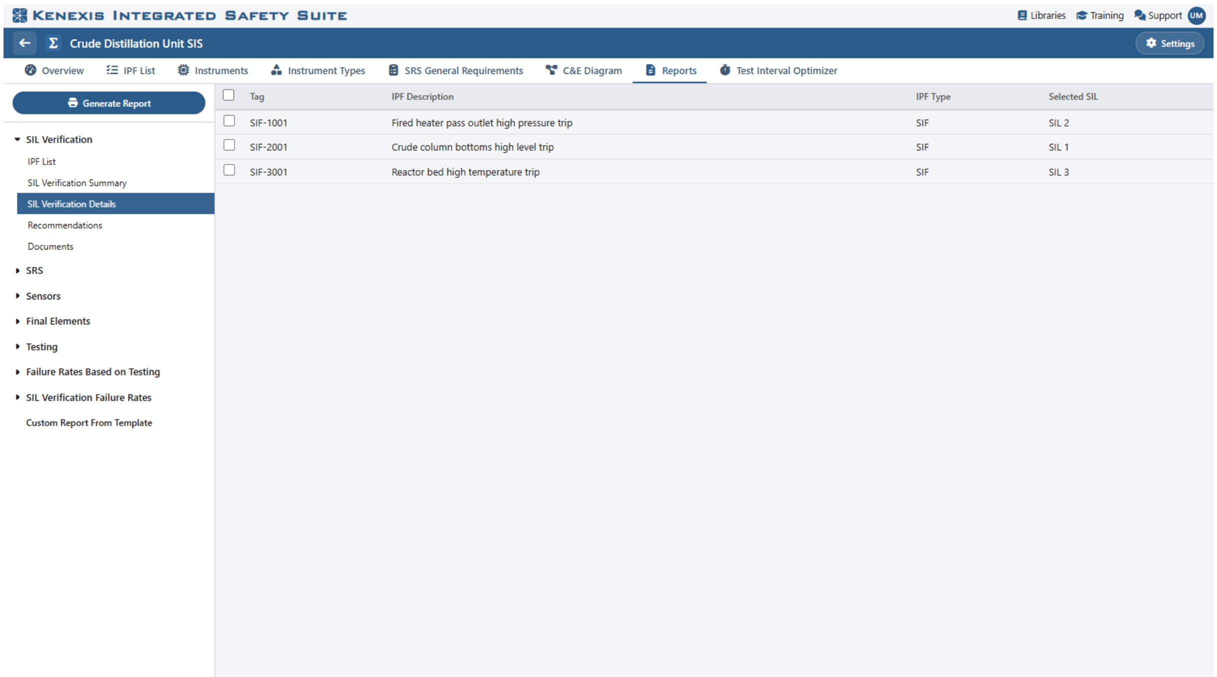

Figure 13-1. The Reports tab. Categories expand to individual reports; the selection grid chooses which items the report covers.

Table 13-1. Report categories, the reports in each, and their output formats.

Category	Reports	Format
SIL Verification	IPF List; SIL Verification Summary; SIL Verification Details (the full appendix-grade calculation record per function); Recommendations; Documents	Word
SRS	SRS General Requirements; IPF Datasheets; Sensor Datasheets; Logic Solver Datasheets; Final Element Datasheets; Cause & Effect Diagrams	Word (C&E in Excel)
Sensors	Setpoint List	Word
Final Elements	Activation Time List	Word
Testing	Current status per category; full test history per category	Word
Failure Rates Based on Testing	Actual failure statistics per type category, computed from the recorded tests	Word
SIL Verification Failure Rates	The failure-rate data used in the calculations, per type category	Word
Custom Report From Template	Your own Word template, populated with study data	Word

Datasheet reports respect the study’s SRS Tracked Fields, so the printed datasheets carry the same field set as the forms. Custom Report From Template accepts a Word (.docx) template of up to 10 MB with merge fields and returns it filled with the study’s data — deliverables in your organization’s own format.

14 Study Settings

The gear button in the workspace header opens the Study Settings dialog — four tabs of per-study configuration.

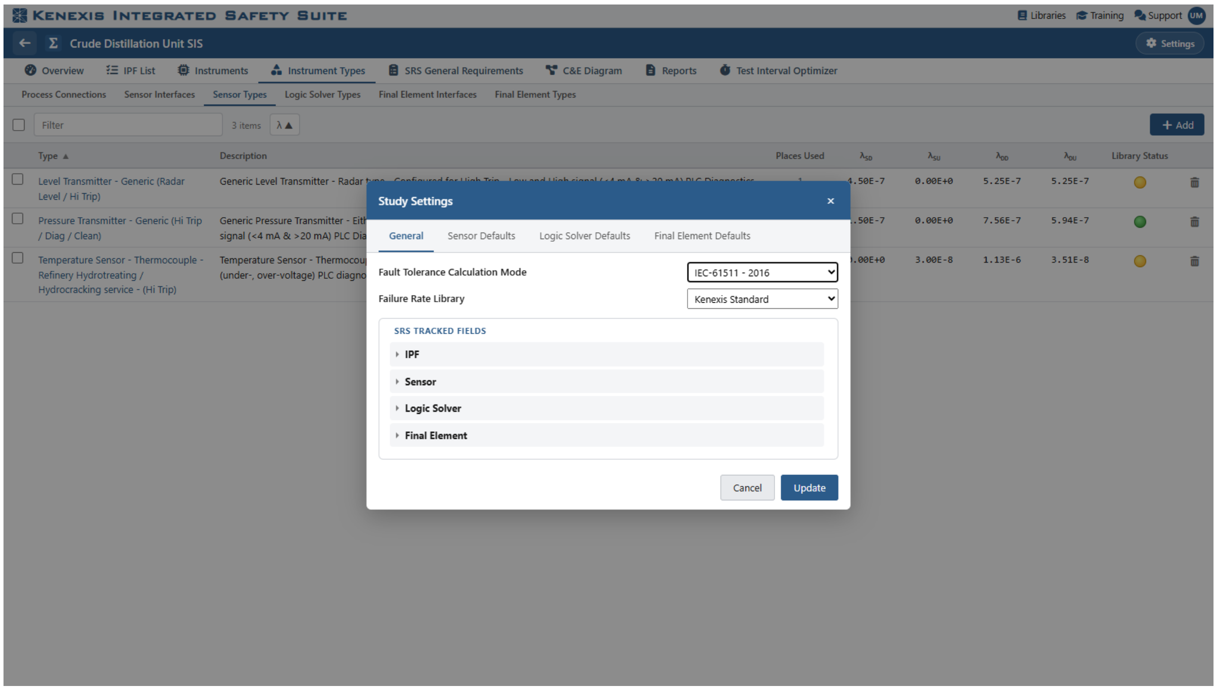

Figure 14-1. Study Settings. The General tab holds the fault-tolerance calculation mode, the failure-rate library, and the SRS tracked-field tree; the remaining tabs set new-instrument defaults per category.

14.1 General

- **Fault Tolerance Calculation Mode** — IEC-61511 - 2003 or IEC-61511 - 2016. This selects which edition's minimum hardware fault tolerance requirements the verification applies (section 16.9); changing it recalculates the whole study.
- **Failure Rate Library** — the library that feeds the instrument-type picker (section 6.3), chosen from the libraries you are authorized to use.
- **SRS Tracked Fields** — the checkbox tree that selects which SRS datasheet fields appear on the IPF, sensor, logic solver, and final element forms and reports (section 8.2).

14.2 Instrument Defaults

The Sensor Defaults, Logic Solver Defaults, and Final Element Defaults tabs seed every newly created instrument of that category: the test interval, the MRT, and the full advanced-parameter set (the five options with their beta factor, diagnostic interval, test duration, proof test coverage, useful life, and comparison settings). Set them once to match your modeling conventions, and new instruments start out consistent.

15 File Attachments

IPFs and instruments each carry a Files tab for supporting documents — vendor datasheets, sizing calculations, photographs, test records. Drop files onto the tab or click to browse; files upload immediately as they are selected, up to twenty files per selection and 50 MB per file. Each entry lists the file's name, size, uploader, and timestamp, with a download link and a delete action. Attachments follow their parent: they are available wherever the IPF or instrument dialog is opened.

16 The Calculations Behind Kenexis SIS

This chapter documents the calculation engine: the equations Kenexis SIS evaluates, when each term applies, and how instrument results roll up to the function verdict. The worked examples use round illustrative numbers.

16.1 Conventions and Units

- **Failure rates** — all λ values are failures per hour. FIT entry divides by 10^9 on the way in.
- **Test interval** — entered in months and converted to hours: τ = TI ÷ 12 × 8,760.
- **MRT, diagnostic interval, test duration, comparison interval** — hours. Useful life — years, converted at 8,760 hours per year.
- **Results** — PFDavg is dimensionless; STR is trips per hour; MTTF-S is years; RRF = 1 ÷ PFDavg.

Calculations run automatically: saving an instrument recalculates it and every function that uses it; editing an instrument type recalculates every instrument of that type; importing a study or changing the fault-tolerance mode recalculates the whole study. The instrument dialog additionally previews the same equations live as you edit.

16.2 Failure-Rate Decomposition

Every instrument type resolves to four failure-rate splits — safe or dangerous, detected or undetected. When data is entered in the summary form (failure rate λ, percent safe, and the two diagnostic coverages), the splits are:

$$\lambda_{\mathrm{SD}} = \lambda \cdot s \cdot c_S$$
$$\lambda_{\mathrm{SU}} = \lambda \cdot s \cdot (1 - c_S)$$
$$\lambda_{\mathrm{DD}} = \lambda \cdot (1 - s) \cdot c_D$$
$$\lambda_{\mathrm{DU}} = \lambda \cdot (1 - s) \cdot (1 - c_D)$$

s = percent safe ÷ 100; c_S = safe coverage ÷ 100; c_D = dangerous coverage ÷ 100.

An instrument then sums the splits of everything in its subsystem: the device type plus, for sensors, the process connection and up to two sensor interfaces, and for final elements up to two final element interfaces. These totals — the ones shown in the dialog's Subsystem Failure Rates table — are the λ values used below.

16.3 Voting Combinatorics

For a group of N identical devices voting M-out-of-N, the calculation uses three derived constants. K is the number of dangerous failures that defeats the function — the function survives until only M−1 healthy devices remain:

$$K = N - M + 1$$

$$k_n = \frac{N!}{K! \cdot (N-K)!}$$

$$k_s = \frac{N!}{(N-M)!}$$

k_n counts the combinations of K dangerous failures; k_s counts the ordered ways M devices can fail spuriously.

Two building blocks recur. Failures revealed only by proof testing average over the test interval ("after logic"); failures repaired as they occur average over their repair window ("before logic"):

$$\mathrm{PFD}_{\mathrm{after}} = \frac{k_n}{K+1} \cdot (\lambda \cdot \tau)^K$$

$$\mathrm{PFD}_{\mathrm{before}} = k_n \cdot \left(\lambda \cdot \left(\mathrm{MTTR} + \frac{T_{\mathrm{int}}}{2}\right)\right)^K$$

τ = test interval; T_int = the applicable repair-cycle interval; MTTR = mean time to restore.

16.4 PFD Contributions

The instrument's PFDavg is the sum of six contributions; each applies only when its option is active.

Dangerous Undetected (DU)

The core term — dormant dangerous failures found at proof test:

$$\mathrm{PFD}_{\mathrm{DU}} = \frac{k_n}{K+1} \cdot (\lambda_{\mathrm{DU}} \cdot \tau)^K$$

Worked example

A single transmitter (1oo1, so K = 1, k_n = 1) with $\lambda DU = 2\times10^{-6}$ per hour tested every 12 months (τ = 8,760 h): PFD = $\lambda DU \cdot \tau \div 2 = 8.76\times10^{-3}$. As a 1oo2 pair (K = 2) the same data gives $(\lambda DU \cdot \tau)^2 \div 3 = 1.02\times10^{-4}$ — redundancy buys nearly two orders of magnitude before common cause is considered.

Detected Failures (DD)

With Trip on Detected Failure switched on, detected failures vote to trip and contribute nothing to PFD (they appear in the STR instead, section 16.5). Switched off, detected failures are alarmed and repaired, and both detected splits contribute across the diagnostic interval DI:

$$\mathrm{PFD}_{\mathrm{DD}} = k_n \cdot \left((\lambda_{\mathrm{SD}} + \lambda_{\mathrm{DD}}) \cdot \left(\mathrm{MTTR} + \frac{\mathrm{DI}}{2}\right)\right)^K$$

Common Cause (CC)

With Common Cause on and redundant voting (N ≠ M), a beta fraction of the dangerous undetected rate strikes all channels at once and behaves like a single 1oo1 component:

$$\mathrm{PFD_{CC}} = \beta \cdot \lambda_{\mathrm{DU}} \cdot \frac{\tau}{2} + \beta \cdot \lambda_{\mathrm{DN}} \cdot \frac{\mathrm{UL}}{2}$$

β = beta factor; UL = useful life in hours; the second term applies with imperfect testing (below).

Worked example

The 1oo2 pair above with β = 10%: PFD_CC = 0.1 × 8.76×10⁻³ = 8.76×10⁻⁴ — nearly ten times the pair's independent contribution. This is why the beta factor deserves as much attention as the failure rates.

External Comparison (Compare)

With External Comparison on, the comparison reveals a further share of dangerous failures beyond the internal diagnostics. The extra coverage is repaired on the comparison interval ECI:

$$c_{\mathrm{extra}} = \max\left(0, \mathrm{ECC} - \frac{\lambda_{\mathrm{DD}}}{\lambda_{\mathrm{D,tot}}}\right)$$

$$\mathrm{PFD_{cmp}} = k_n \cdot \left(\lambda_{\mathrm{D,tot}} \cdot c_{\mathrm{extra}} \cdot \left(\mathrm{MTTR} + \frac{\mathrm{ECI}}{2}\right)\right)^K$$

ECC = comparison coverage; λ_D,tot = λ_DD + λ_DU (+ λ_DN); the DU term is reduced by the same extra coverage.

Online Testing (Test)

With Online Testing on, the device is bypassed for the test duration TD each interval, and that bypass time is unavailability:

$$\mathrm{PFD_{test}} = \frac{\mathrm{TD}}{\tau}$$

Never-Detected Failures (DN)

With Imperfect Testing on, only the proof-test-coverage fraction of dangerous undetected failures is found at test; the remainder persists until the device is replaced at the end of its useful life:

$$\lambda_{\mathrm{DN}} = \lambda_{\mathrm{DU}} \cdot (1 - \mathrm{PTC})$$

$$\lambda_{\mathrm{DU}}^{\dagger} = \lambda_{\mathrm{DU}} \cdot \mathrm{PTC}$$

$$\mathrm{PFD_{DN}} = \frac{k_n}{K+1} \cdot (\lambda_{\mathrm{DN}} \cdot \mathrm{UL})^K$$

PTC = proof test coverage; λ_DU† replaces λ_DU in the DU term; UL = useful life in hours.

Total

$$\mathrm{PFD_{avg}} = \mathrm{PFD_{DU}} + \mathrm{PFD_{DD}} + \mathrm{PFD_{CC}} + \mathrm{PFD_{cmp}} + \mathrm{PFD_{test}} + \mathrm{PFD_{DN}}$$

16.5 Spurious Trip Rate

The STR counts trips per hour caused by instrument failures rather than process demands. A spurious trip needs M coincident spurious channel failures within the repair window:

$$\mathrm{STR_{SU}} = k_s \cdot \lambda_{\mathrm{SU}}^{M} \cdot \mathrm{MTTR}^{M-1}$$
$$\mathrm{STR_{DD}} = k_s \cdot (\lambda_{\mathrm{SD}} + \lambda_{\mathrm{DD}})^{M} \cdot \mathrm{MTTR}^{M-1}$$
$$\mathrm{STR_{CC}} = \beta \cdot \lambda_{\mathrm{SU}}\big(+\beta \cdot (\lambda_{\mathrm{SD}} + \lambda_{\mathrm{DD}})\big)$$

STR_DD applies when Trip on Detected Failure is on; STR_CC applies with common cause and M > 1, its second term when trip-on-detected is also on.

$$\mathrm{STR} = \mathrm{STR_{SU}} + \mathrm{STR_{DD}} + \mathrm{STR_{CC}}$$

For a 1ooN group, M = 1: any single spurious failure trips the group, so STR_SU = N·λSU and the MTTR window and common-cause term drop out.

16.6 Safe Failure Fraction

The SFF measures how much of the total failure rate is either safe or detected — everything except the dangerous failures that stay hidden:

$$\mathrm{SFF} = 1 - \frac{\lambda_{\mathrm{DU}} + \lambda_{\mathrm{DN}}}{\lambda_{\mathrm{DU}} + \lambda_{\mathrm{DN}} + \lambda_{\mathrm{DD}} + \lambda_{\mathrm{SD}} + \lambda_{\mathrm{SU}}}$$

The SFF appears in the instrument dialog and drives the architectural constraint assessment of section 16.9.

16.7 Subsystems and Group Logic

A function's sensor and final element subsystems can hold several voted groups; the group logic (section 4.4) sets how their results combine:

Table 16-1. Subsystem roll-up by group logic.

Group logic	Subsystem PFD	Subsystem STR
1oo1 (single group)	$\mathrm{PFD} = \sum \mathrm{PFD}_i$	$\mathrm{STR} = \sum \mathrm{STR}_i$
1ooX (any group suffices)	$\mathrm{PFD} = \prod \mathrm{PFD}_i$	$\mathrm{STR} = \sum \mathrm{STR}_i$
XooX (all groups required)	$\mathrm{PFD} = 1 - \prod (1 - \mathrm{PFD}_i)$	$\mathrm{STR} = N \cdot (\prod \mathrm{STR}_i) \cdot 72^{N-1}$

With XooX logic, redundant groups back each other up against spurious trips: all N groups must trip spuriously within a coincidence window (72 hours) for the function to trip. The logic-solver subsystem always sums its members' contributions.

16.8 Function Results

With all three subsystems calculated, the function-level results follow:

$$\mathrm{PFD_{IPF}} = \mathrm{PFD_{sensors}} + \mathrm{PFD_{logic}} + \mathrm{PFD_{final}}$$

$$\mathrm{RRF} = \frac{1}{\mathrm{PFD_{IPF}}}$$

$$\mathrm{STR_{IPF}} = \mathrm{STR_{sensors}} + \mathrm{STR_{logic}} + \mathrm{STR_{final}}$$

$$\frac{1}{\mathrm{MTTFS}} = \frac{1}{\mathrm{MTTFS_{sensors}}} + \frac{1}{\mathrm{MTTFS_{logic}}} + \frac{1}{\mathrm{MTTFS_{final}}}$$

Each subsystem's MTTFS is 1 ÷ (STR × 8,760) years.

The achieved SIL comes from the standard low-demand bands:

Table 16-2. Achieved SIL bands for low-demand mode.

Overall PFDavg	Achieved SIL
PFD ≤ 1×10^{-4}	SIL 4
1×10^{-4} < PFD ≤ 1×10^{-3}	SIL 3
1×10^{-3} < PFD ≤ 1×10^{-2}	SIL 2
1×10^{-2} < PFD ≤ 1×10^{-1}	SIL 1
PFD > 1×10^{-1}	N/A

Max SIL Approved combines the PFD band with the fault-tolerance assessment: the function is approved to a SIL only when both its PFD and every subsystem's fault-tolerance capability support it. Minimum Fault Tolerance Achieved compares that capability with the selected SIL and reads Yes or No — the second half of the verification verdict.

16.9 Hardware Fault Tolerance

IEC 61511 imposes minimum hardware fault tolerance (HFT) — the number of dangerous failures a subsystem must tolerate while continuing to function — independent of the calculated PFD. Each voted group achieves HFT = N − M; a 1ooX subsystem adds its groups' tolerances (plus one for the redundancy between groups), while 1oo1 and XooX subsystems take their weakest member.

The achieved tolerance is then weighed against the requirement using the study's Fault Tolerance Calculation Mode:

- **IEC-61511 - 2003 mode** — applies the first edition's Tables 5 and 6. The SFF-based architectural constraints of IEC 61508-2 (Tables 2 and 3, by architecture type A or B) are always evaluated as the alternative route, and the more favorable of the two governs. A device qualified by prior use (the Reduce Fault Tolerance credit on its SRS datasheet) relaxes the requirement by one unit of tolerance, as the 2003 edition permits.
- **IEC-61511 - 2016 mode** — applies the second edition's Table 6: HFT ≥ 1 supports up to SIL 3, HFT 0 up to SIL 2, with the high-demand SIL 2 case requiring one more unit of tolerance. The IEC 61508-2 route remains available, and again the more favorable result governs.

The SFF-based capability adjustment per IEC 61508-2 is:

Table 16-3. SFF-based capability adjustment per IEC 61508-2.

SFF	Type A adjustment	Type B adjustment
Below 60%	capability = HFT	capability = HFT − 1
60% to below 90%	capability = HFT + 1	capability = HFT
90% to below 99%	capability = HFT + 2	capability = HFT + 1
99% and above	capability = HFT + 2	capability = HFT + 2

A SIL-certified logic solver type carries its certification directly: its capability is the certified SIL. Black Box Model types carry the hardware fault tolerance and maximum SIL you assign to them.

16.10 Continuous Mode

A function whose Mode of Operation is Continuous is judged on its dangerous failure frequency PFH (per hour) rather than PFDavg, following IEC 62061. Each instrument's PFH combines its effective dangerous rate with the channel-equivalent downtime; subsystems combine by the same group-logic rules; and the achieved SIL uses the PFH bands (SIL 2 below 10^{-6} per hour, SIL 3 below 10^{-7}, and so on), together with an SFF-and-HFT architectural constraint per subsystem. The dialog relabels its results accordingly.

16.11 Black Box Model Values

A BBM instrument bypasses the equations above: its PFDavg and STR come straight from the type record — typed in or pushed from a linked Kenexis Fault Tree fault tree (section 6.6) — and its fault-tolerance data from the type's HFT and maximum SIL fields. Everything downstream (subsystem roll-up, function results, SIL banding) treats those values exactly like calculated ones.

16.12 Optimizer Algorithm

The Test Interval Optimizer searches downward from the maximum interval. A candidate interval is applied to the devices being optimized, the full engine recalculates, and each selected function passes when its overall PFD is at or below 1 ÷ Required RRF. In study-wide mode a single common interval steps down until every selected function passes or the minimum is reached. In per-IPF mode each function steps its own devices down independently, and a device shared between functions keeps the shortest interval any of them needs. Because every step uses the real calculation engine, accepted results always match what the study will show.

16.13 Beta Factor Estimator Scoring

Each checked defense in the estimator contributes points; the total maps to a beta factor:

Table 16-4. Beta factor by estimator score band.

Score	Field devices	Programmable logic solvers
Below 45	10%	5%
45 to below 70	5%	2%

Score	Field devices	Programmable logic solvers
70 to below 120	2%	1%
120 and above	1%	0.5%

Better-defended installations earn lower beta factors, and programmable logic solvers — with their systematic diagnostics — earn lower values than field devices at every score band.

16.14 Status Indicators

- **Test due date** — the last recorded test date (or the commission date before any tests) plus the test interval. Yellow starts 182 days before the due date; red starts at it (section 10.1).
- **Activation rate** — valid events divided by the operational time from commissioning to now (or to decommissioning). Green while the rate is below the expected demand rate; yellow at or above it (section 11.2).
- **SIL verification light** — green when the floored achieved RRF meets the required RRF and the minimum fault tolerance is satisfied; red otherwise (section 4.2).

16.15 Data Validation

A calculation runs only when its inputs are complete: the instrument needs a type with full failure-rate data, a test interval, and voting, and each active option needs its own inputs (beta factor; diagnostic interval and MRT; test duration; proof test coverage and useful life; comparison coverage and interval). While anything is missing, the result fields stay empty and the study's calculation log names the missing item, so an absent result always has a stated reason.

Appendix A Definitions

Table A-1. Definitions of the terms used in this manual.

Term	Definition
Beta factor (β)	The fraction of a redundant group's failures attributable to a common cause that strikes all channels at once.
Dangerous coverage	The fraction of dangerous failures the device's diagnostics detect.
Deenergize-to-trip (DTT)	Outputs are energized in normal operation; removing power causes the trip.
Demand	A process condition that requires the safety function to act.
Energize-to-trip (ETT)	Outputs are de-energized in normal operation; applying power causes the trip.
Final element	The device (typically a valve) that carries out the safety action.
Hardware fault tolerance (HFT)	The number of dangerous failures a subsystem can sustain and still perform its function.
IPF	Instrumented Protective Function — an instrumented safeguard comprising sensor, logic solver, and final element subsystems.
Logic solver	The system (typically a safety PLC) that evaluates trip logic between sensors and final elements.
MooN voting	M-out-of-N: M of the N installed devices must vote to trip for the group to act.
MRT / MTTR	Mean repair time — the average time from detecting a failure to restoring the device to service.
MTTF-S	Mean time to fail spurious — the average time between spurious trips; the inverse of the STR.
PFDavg	Average probability of failure on demand — the probability the function cannot act when a demand occurs, averaged over the test interval.
PFH	Average frequency of dangerous failure per hour — the high-demand / continuous-mode counterpart of PFDavg.
Proof test coverage	The fraction of dangerous undetected failures a proof test actually reveals.
RRF	Risk reduction factor — 1 ÷ PFDavg.
Safe failure fraction (SFF)	The fraction of the total failure rate that is safe or detected.
SIF	Safety Instrumented Function — an IPF with a SIL requirement.
SIL	Safety Integrity Level — the order-of-magnitude performance band (1–4) defined by IEC 61511 / IEC 61508.
SIS	Safety Instrumented System — the implementation of one or more SIFs.
SRS	Safety Requirements Specification — the documented requirements for the SIS.
STR	Spurious trip rate — the frequency of trips caused by instrument failures rather than process demands.
Subsystem	One of the three parts of an IPF: sensors, logic solvers, or final elements.
λSD, λSU, λDD, λDU	The four failure-rate splits: safe/dangerous × detected/undetected, per hour.
λDN	Dangerous never-detected — the failure rate share that imperfect proof testing can never reveal.

Appendix B Formula Quick Reference

Table B-1. Formula quick reference.

Quantity	Formula
Rate splits	$\lambda_{\mathrm{DU}} = \lambda \cdot (1-s) \cdot (1-c_D), \qquad \lambda_{\mathrm{DD}} = \lambda \cdot (1-s) \cdot c_D$
Voting constants	$K = N - M + 1, \qquad k_n = \frac{N!}{K!\,(N-K)!}, \qquad k_s = \frac{N!}{(N-M)!}$
DU contribution	$\mathrm{PFD_{DU}} = \frac{k_n}{K+1} \cdot (\lambda_{\mathrm{DU}} \cdot \tau)^K$
DD contribution (no trip)	$\mathrm{PFD_{DD}} = k_n \cdot \left((\lambda_{\mathrm{SD}} + \lambda_{\mathrm{DD}}) \cdot \left(\mathrm{MTTR} + \frac{\mathrm{DI}}{2} \right) \right)^K$
Common cause	$\mathrm{PFD_{CC}} = \beta \cdot \lambda_{\mathrm{DU}} \cdot \frac{\tau}{2} + \beta \cdot \lambda_{\mathrm{DN}} \cdot \frac{\mathrm{UL}}{2}$
Online testing	$\mathrm{PFD_{test}} = \frac{\mathrm{TD}}{\tau}$
Never-detected	$\lambda_{\mathrm{DN}} = \lambda_{\mathrm{DU}} \cdot (1 - \mathrm{PTC}), \qquad \mathrm{PFD_{DN}} = \frac{k_n}{K+1} \cdot (\lambda_{\mathrm{DN}} \cdot \mathrm{UL})^K$
STR (spurious)	$\mathrm{STR_{SU}} = k_s \cdot \lambda_{\mathrm{SU}}^M \cdot \mathrm{MTTR}^{M-1}$
SFF	$\mathrm{SFF} = 1 - \frac{\lambda_{\mathrm{DU}} + \lambda_{\mathrm{DN}}}{\lambda_{\mathrm{tot}}}$
Subsystem XooX	$\mathrm{PFD} = 1 - \prod(1 - \mathrm{PFD}_i)$
Function PFD	$\mathrm{PFD_{IPF}} = \mathrm{PFD}_S + \mathrm{PFD_{LS}} + \mathrm{PFD_{FE}}$
RRF	$\mathrm{RRF} = \frac{1}{\mathrm{PFD_{IPF}}}$

End of manual.