IndexForge

WHITE PAPER

FOR BANKS AND CREDIT UNIONS

Every document accounted for.

How IndexForge turns messy ECM exports into clean, verified import files for mergers, acquisitions and ECM migrations, without a single byte of member or customer data leaving your network.

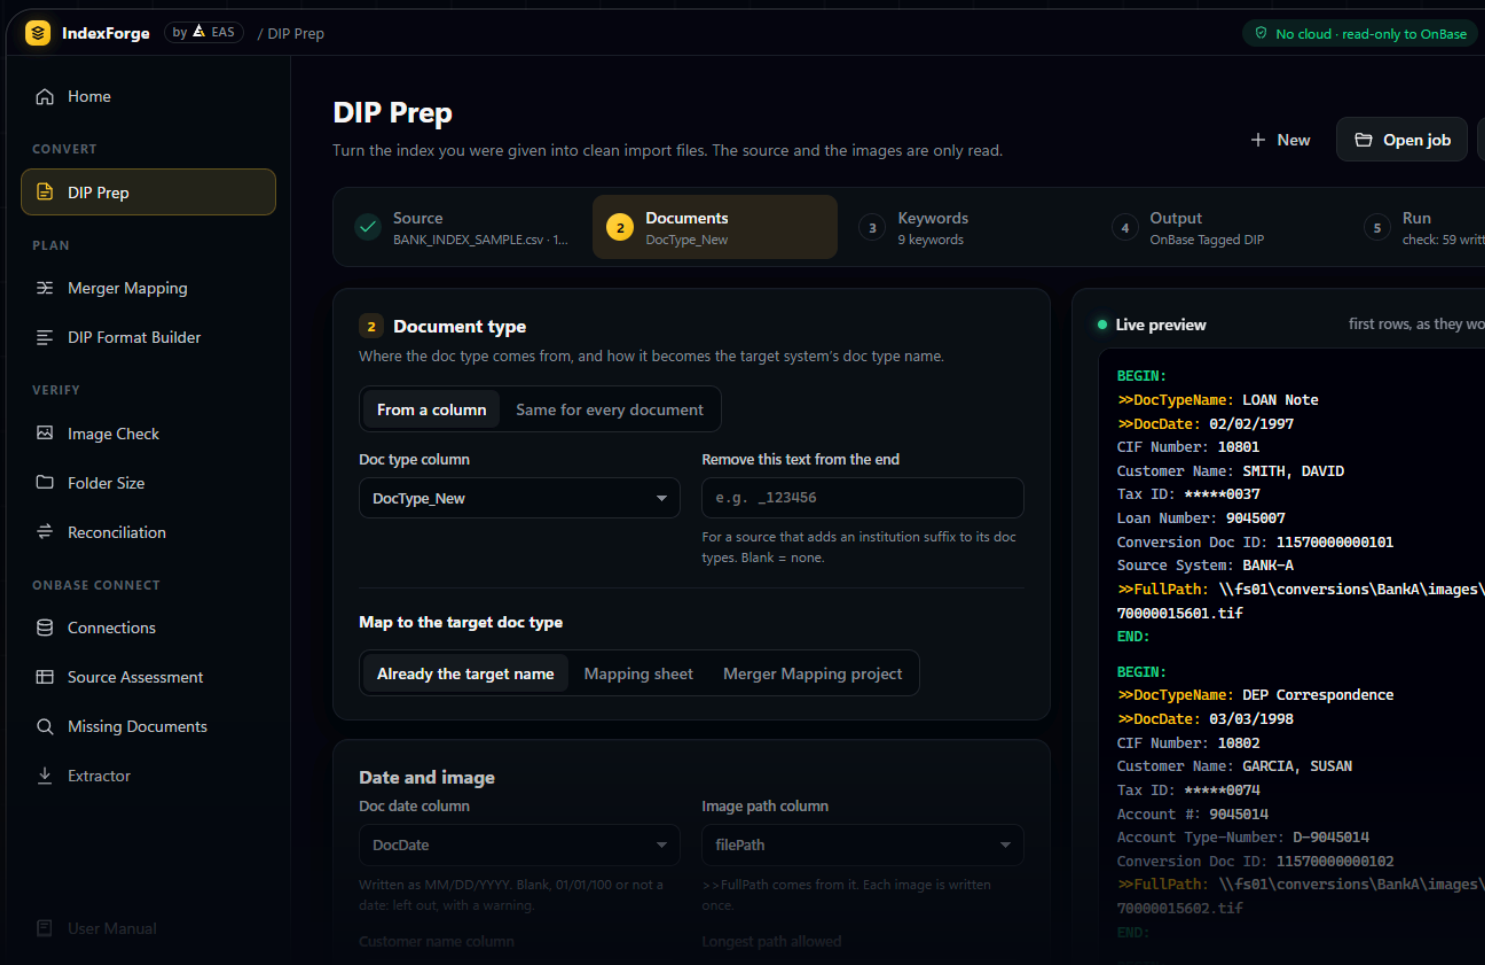

Document conversion · Deconversion · OnBase DIP · Any ECM

Elevate Automation Solutions LLC (EAS) · November 2026

EXECUTIVE SUMMARY

Conversions fail in the index file, not in the ECM.

When institutions merge, change vendors or consolidate charters, millions of signature cards, loan files, statements and correspondence move from one content system to another. The documents are rarely the problem. The **index file** that describes them almost always is.

Exports arrive with names split across columns, several accounts crammed into one row, the literal word NULL where a value should be, images that no longer exist and doc types that mean nothing to the receiving institution. Each of these turns into a document that lands under the wrong account, a keyword that silently shifts, or a file that never arrives at all, and each one surfaces months later in front of a member, a customer or an examiner.

IndexForge is a desktop workbench built for exactly this step. It reads whatever the source system produced, applies institution-grade rules, and writes clean import files for Hyland OnBase or for any other ECM. Every source row ends up in one of two places: **written**, or **held back with a plain-English reason**. Nothing slips through unexplained.

100%
of source rows accounted for: written + held always equals rows read

~34,600
rows per second in internal testing (200,000-row bank index)

0
cloud services: data never leaves your network

5 → 4
source types in, output formats out, through one common document model

Who this paper is for

- **Conversion and merger integration leads** who own the timeline and the go-live weekend.
- **ECM and OnBase administrators** who build the DIP processes and answer for what lands.
- **Deposit and loan operations** who need every signature card and note under the right account.
- **Information security, risk and compliance** who must know where sensitive data goes, and prove it.

> The goal is not a faster spreadsheet. It is a conversion you can explain row by row, to a member, an auditor or your board.
>
> Every module that follows serves that goal.

THE CHALLENGE

Six ways a document conversion goes quietly wrong

None of these is exotic. Every conversion team has seen all six, usually in the same file.

Broken index files

"SMITH, JOHN" without quotes becomes two columns. Three accounts share one row. SQL exports drop the header line. Every row after the first bad comma is shifted.

Doc types that do not match

Bank A's "DEP - Sig Card" is Bank B's "DEP Signature Cards". Hundreds of these are matched by hand in spreadsheets that are emailed back and forth.

Missing and duplicate images

Files deleted outside the ECM, a second or third disk-group copy nobody knew about, and two rows pointing to the same image. Each one is a document that never arrives, or arrives twice.

Keywords that shift

Multi-instance keyword groups (owner, address, city) pair up by order. One blank value and every later owner is attached to the wrong address, with no error at all.

Sensitive data everywhere

Index files carry names, Tax IDs and account numbers. Cleaning them in spreadsheets, scripts and email attachments spreads that data far beyond where it belongs.

No proof of completeness

After go-live, someone asks "did everything come over?" Without a row-by-row account and a reconciliation, the honest answer is "we think so".

A conversion is only as good as its index file. IndexForge treats that file as the thing to be engineered, tested and proven.

The design principle behind every module in this paper.

HOW IT WORKS

One pipeline, any source, any target

Every source is read into the same vendor-neutral document model: a document type, a date, keywords that may hold several values, and the file itself. Rules run once against that model, and writers turn it into whatever the receiving system imports.

1 — Ingest

CSV, tab or pipe files, XML exports, fixed-width records, folder trees, or existing DIP files. Huge files are streamed, never loaded whole.

2 — Map

Source doc types become target doc types through a mapping sheet or a Merger Mapping project. Keywords come from columns, fixed values, templates or a running sequence.

3 — Validate

Repairs what can be safely repaired. Holds back what cannot, with the reason. Checks every image exists and is used once.

4 — Write

OnBase Tagged or Ordered DIP, or a generic CSV or XML index for any ECM. Split by document count or by MB of images.

5 — Verify

HTML report, held and warning lists, a SHA-256 manifest, image checks and reconciliation against the target.

Built for the go-live weekend

- **Test** writes the first 25 documents to a file you can open and eyeball.
- **Check** reads the whole source and reports exactly what would be written, held and split, without creating a single output file.
- **Create** writes the import files, never overwriting anything, and warns before the same source file is created twice.
- A **live preview** shows the first documents exactly as they will be written while you configure, so mistakes are visible in seconds, not after ingestion.

Inputs and outputs

READS	WRITES
Delimited text (auto-detects comma, tab, pipe, semicolon; UTF-8 or ANSI)	OnBase Tagged DIP
XML exports (record element found automatically)	OnBase Ordered DIP
Fixed-width / mainframe records	Generic CSV index
Folder structures, with names parsed by pattern	Generic XML manifest
Existing tagged DIP index files	More writers can be added

FEATURE · DIP PREP

From the file you were given to the file OnBase expects

A five-step guided workflow (Source, Documents, Keywords, Output, Run) with the result previewed live beside every setting.

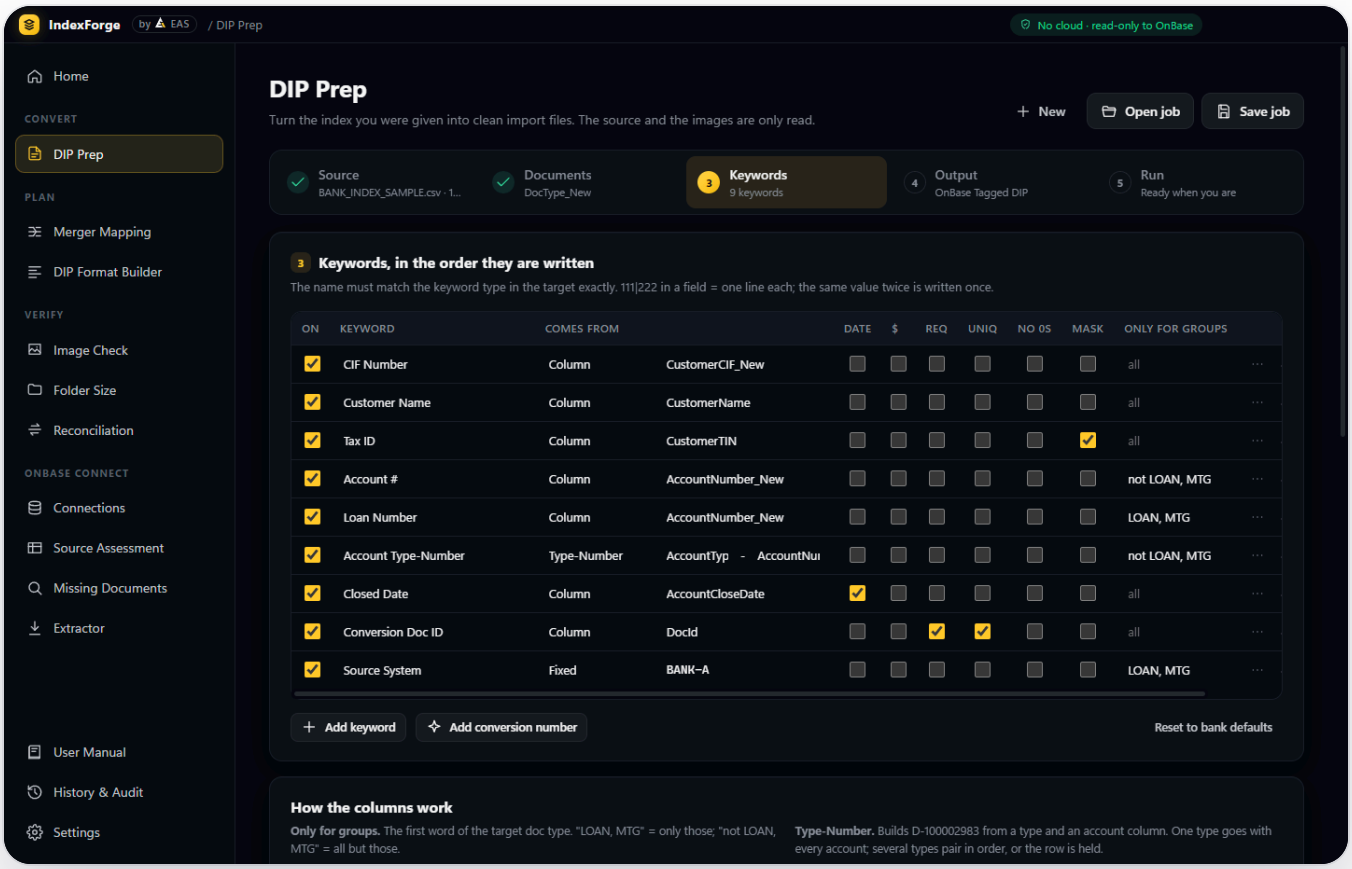

The keyword grid: each keyword's source, date and currency formatting, Required and Unique checks, leading-zero removal, masking, and doc-type group rules. Shown with fictitious demo data.

Keywords your way

From a column, a fixed value, a template such as **{Branch}-{Account}**, a type-and-number pair (**D-9045335**), or a running **conversion number** stamped on every document for reconciliation.

Rules by document group

Loan documents get Loan Number; deposit documents get Account # and account type. One grid, driven by the first word of the target doc type.

Transforms built in

Case, find-and-replace, padding, prefixes, maximum length, MM/DD/YYYY dates and OnBase currency format, without a separate script.

DATA QUALITY

Repair what is safe. Hold what is not. Never guess.

The rules come from real bank conversion files. A row that cannot be rebuilt with certainty is held, never best-guessed into the wrong account.

SITUATION IN THE SOURCE	WHAT INDEXFORGE DOES
Names with unquoted commas (SMITH, JOHN)	Puts the name back together when the rest of the row still lines up, and lists it as a warning. Otherwise HELD
Several accounts, names or Tax IDs on one row	Recognizes the bank's 14-column layout, rebuilds the row, writes one keyword line per value and drops exact repeats
NULL, NULL-NULL or "-" as a value	Treated as empty, never written as text
Doc type missing, unmapped, or mapped to two targets	HELD with the doc type named, so the mapping can be fixed
Image missing, path over 260 characters, or an image used twice	HELD; a repeated image is called out in capitals before any file is created
Required keyword blank, or a Unique keyword repeated	HELD with the earlier row number
Placeholder dates (01/01/100) or invalid dates	Left out, with a WARNING; valid dates are rewritten as MM/DD/YYYY
Account types that do not pair with account numbers	One type is applied to every account; several types pair in order, or the row is HELD
Uneven multi-instance keyword groups	WARNING before ingestion, because OnBase DIP would otherwise attach values to the wrong instance
Images moved to a new server or share	Path rewrite: the start of every path is replaced before it is checked and written
Very large batches	A new import file after N documents or N MB of images. A document is never split across files.

The Held file is a work list

Every held row is written to a CSV with its reason and every original column. Fix those rows and run them again as their own source.

Fast enough for the largest files

In internal testing on a standard Windows workstation, a 200,000-row bank index was processed in under 6 seconds (image checks off). Image checks over a network share run many lookups in parallel.

FEATURE · MERGER MAPPING

Map two institutions in hours, not weeks

Load Bank A's and Bank B's doc types and keywords, side by side on one sheet or in two files. IndexForge scores every possible match, understanding banking shorthand such as Acct, Sig, LN, Corresp and Agmt.

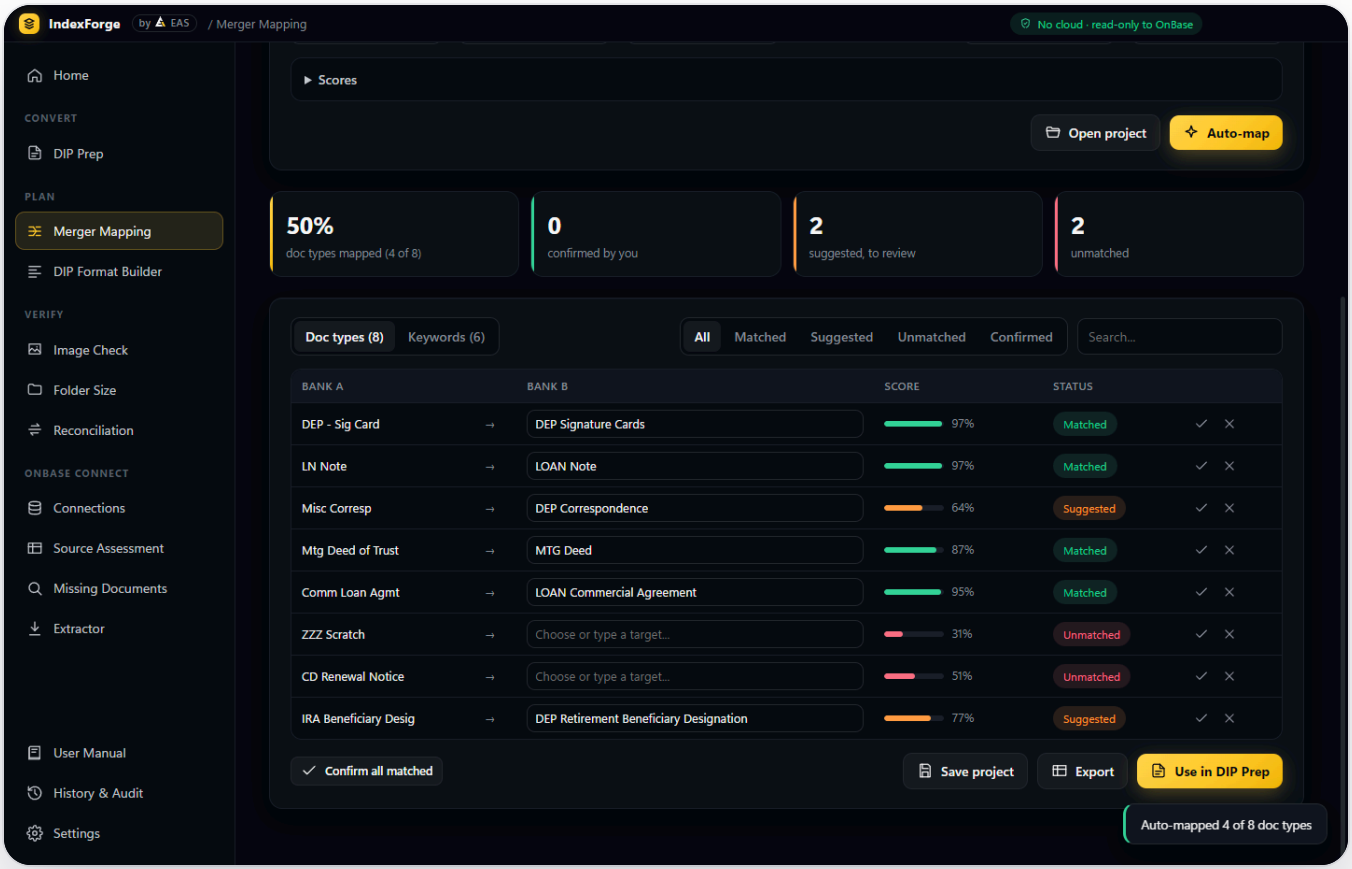

Scored matches with status: Matched, Suggested for review, Unmatched or Confirmed. Targets are chosen from a searchable list of the receiving institution's names.

No more spreadsheet ping-pong

- Confirm all strong matches in one click; review the rest in place.
- Save the project, reopen it, and export it to Excel for sign-off.
- Use it directly as the doc-type map in DIP Prep: no retyping.
- It is never 100%, and the tool says so. Anything below the score you set stays blank for a person to decide.

FEATURE · DIP FORMAT BUILDER

Let the index file configure the process

Point it at a tagged DIP file and it produces the OnBase DIP process settings, the field order (tag to keyword type), the longest value and most values for each tag, and the keywords each doc type must carry. It flags tags over OnBase's 20-character limit and keywords that need multiple instances, all in a workbook your OnBase administrator can follow step by step.

FEATURES · VERIFICATION

Prove it, before and after go-live

Three tools answer the questions every conversion is eventually asked: is it all there, is it the right size, and did it all arrive?

Image Check

Confirms every file in an index or extract list exists. If it is not at copy 1, copies 2 and 3 are tried. Missing files are listed by doc type before anyone builds a DIP.

Folder Size

Counts files and bytes actually on disk, by extension and folder, against the numbers the source system reports. Gaps usually mean files deleted outside the ECM.

Reconciliation

Extracted against landed, on doc handle or IndexForge's conversion number: missing, extra, duplicated, or matched with a different doc type.

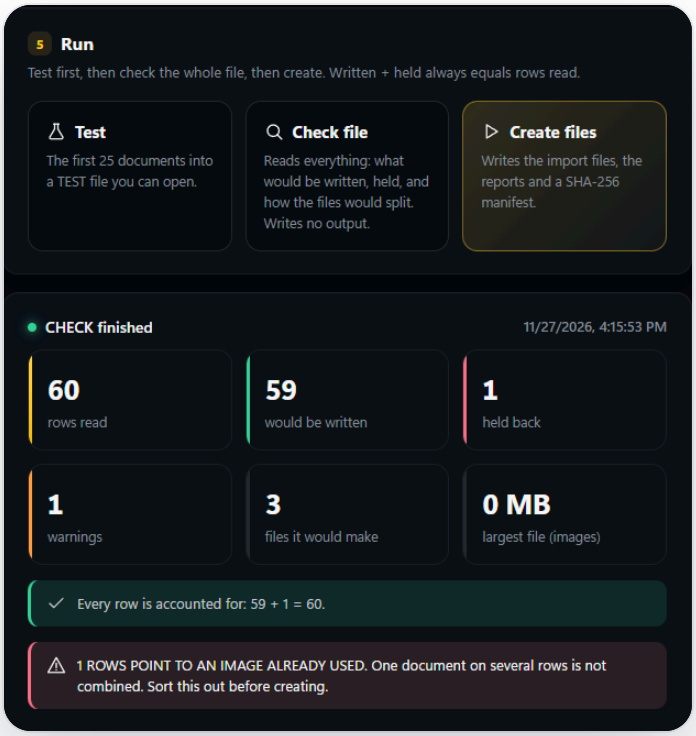

A Check run: every row accounted for, the repeated image flagged before anything is written.

The evidence package

Every Check and Create run leaves a complete, self-explaining record:

- **HTML report**: summary, doc types, keyword fill rates, held rows, warnings and output files.
- **Held CSV**: every held row with its reason, ready to fix and re-run.
- **Warnings CSV**: everything written with a caveat.
- **SHA-256 manifest** of every import file, so you can prove the file ingested is the file produced.
- **Audit log**: who ran what, on which machine, when, with counts.

USE CASES

Where institutions put IndexForge to work

The same engine serves every direction a document can move.

01 Bank merger or acquisition

The acquired bank's documents move into the acquirer's OnBase. Merger Mapping matches doc types and keywords, DIP Prep builds the DIP files, and Reconciliation proves arrival on the conversion number.

Merger Mapping → DIP Prep → Reconciliation

02 Credit union merger

Two charters become one. Member documents keep their account history while doc types and keyword names are standardized to the continuing credit union's taxonomy.

Merger Mapping → DIP Prep (templates, transforms)

03 Migrating into OnBase

Leaving another ECM or an imaging system for OnBase. Its CSV, XML or folder export is read as-is; DIP Format Builder turns the result into a ready-to-configure DIP process.

Any source → DIP Prep → DIP Format Builder

04 Deconversion to another ECM

Moving documents out to a different platform, or handing a book of business to a buyer. Output is a clean, vendor-neutral CSV or XML index the receiving system can import.

DIP / export → DIP Prep → Generic CSV or XML

05 Pre-ingestion audit

Before a single file is imported: run Check on the bank's index, Image Check on the extract, and Folder Size on the share. Fix what is held while there is still time.

Check → Image Check → Folder Size

06 Repair and re-process

A DIP failed, or an old index needs new keywords, a new path or a new split. Read the existing DIP file as a source, re-map and re-write it, with nothing typed by hand.

Existing DIP → DIP Prep → new DIP

> Typical path for a merger: map doc types → Test 25 documents → Check the full file → fix the Held list → Create → ingest → reconcile.

SECURITY AND GOVERNANCE

Designed for data you are not allowed to lose

Index files hold names, Tax IDs and account numbers. IndexForge is built on the assumption that none of it should ever leave the institution, or even the machine running the conversion.

No cloud, by design

No cloud service, no telemetry, no update channel. The only connections are the ones you configure: your OnBase database, read-only, and your file shares.

Sandboxed interface

The window has no direct file or system access. It can only request a fixed, short list of actions, and it opens only files it wrote or that you chose.

Masking

Tax IDs are masked by default (last four only) in previews and reports. Any column or keyword can be masked, including in the Held file.

Read-only sources

Files, images and OnBase are only ever read: fixed SELECT statements only, and a red warning if a login could change data. Outputs are never overwritten.

Integrity

A SHA-256 hash of every import file is recorded at creation, so tampering or a mixed-up file is detectable before ingestion.

Audit trail

User, machine, time, source, output, counts and hashes for every run, kept locally in the user's profile for review.

Supporting your information security program

These controls are designed to support the data-handling expectations institutions apply under the GLBA Safeguards requirements and FFIEC IT guidance: least privilege, data minimization, integrity and accountability. IndexForge supports your own controls and policies; it does not replace them, and your institution's review of the tool remains essential.

STRAIGHT TO THE SOURCE SYSTEM

Connect to OnBase, read-only

IndexForge also connects directly to an OnBase database (SQL Server or Oracle) with a read-only login, so assessment, extraction and reconciliation happen in the same tool. It checks the schema before it runs anything.

AVAILABLE

Source Assessment

Doc type groups, document counts and sizes, status (active, trash), Z and test doc types, disk groups and every doc type's keywords, as the assessment workbook, with date filters such as "last five years".

AVAILABLE

Extractor

Multi-threaded extraction in priority order (signature cards first; negative priorities allowed), latest revision only, resumable, with a gap-free conversion number on every document and an index ready for DIP Prep.

AVAILABLE

Missing documents and live reconciliation

Files the database lists that are gone from disk, across copies 1 to 3; and a list straight from the receiving OnBase to reconcile on doc handle or conversion number.

PLANNED

Encrypted documents (Unity API)

Disk groups encrypted by OnBase will be extracted through the Hyland Unity API, using the institution's own license. DIP and COLD process associations will be added to the assessment.

Database access uses Hyland's read-only viewer account or an equivalent login. Passwords are encrypted by Windows for the signed-in user, every OnBase task runs in its own short-lived process, and every connection and run is in the audit log. Planned items are subject to change.

GETTING STARTED

A low-risk first conversion

1

Pick one group

Start with one doc type group, such as deposit signature cards, and its index file.

2

Test and Check

Eyeball 25 documents, then Check the whole file and work the Held list.

3

Create and ingest

Create the DIP files and import them into a test OnBase environment first.

4

Reconcile

Prove every document arrived, then scale to the full conversion.

IndexForge is a product of Elevate Automation Solutions LLC (EAS). © 2026 Elevate Automation Solutions LLC. All rights reserved. OnBase and Unity are trademarks or registered trademarks of Hyland Software, Inc. All other names are the property of their respective owners. Screenshots show fictitious demo data. Performance figures are from internal testing and vary with hardware, file layout and network.