

Xceptor Confirmations Agent

Read any confirmation.

Extract every field.

Deliver structured data in seconds.

No templates required.

Why template-based confirmation extraction breaks down



Variation in templates slows onboarding

Every counterparty uses a different format, layout, and terminology, turning scale into friction and delaying downstream processing.



Frequent format changes make template-based processes fragile

When formats change, templates break. Operations teams absorb continuous maintenance or fall back to manual re-keying, introducing error risk and delaying settlement.



Accuracy caps increase manual effort

Custom AI models require sample uploads, manual labelling, and iterative retraining – and still plateau at around 60-80% accuracy, leaving large volumes of confirmations still dependent on manual review.

Pre-built extraction for any OTC confirmation format

No templates.

No model training.

No operational drag.



The Xceptor Confirmations Agent enables your team to extract structured trade data from any OTC confirmation document – including PDFs and emails – without building templates or training models.

The Confirmations Agent is a pre-built, domain-trained extraction agent that slots into your existing Xceptor process as a single drag-and-drop action. Using a multi-step agentic workflow, the Agent **extracts, validates, and cross-checks** trade economics before returning a structured and highly accurate output.

Every extracted field comes with a confidence score and a reference back to the source document. Low-confidence outputs route automatically to your existing exception handling workflow.

Key capabilities: Designed for throughput, accuracy, and control



No template configuration

Process any format without counterparty-specific setup, template work, labelling, or retraining.



High 90s extraction accuracy

A multi-step agentic validation loop eliminates the ~80% ceiling of single-prompt LLM approaches. Eight model calls. One validated, structured output.



Built-in trade validation

Extracted fields are checked against defined validation criteria before downstream use, so teams act on data they can trust.



Human oversight by design

Confidence scores and source references accompany every extracted field; exceptions surface in Xceptor's standard exception management workflow with full audit trail.



Native Xceptor integration

Extraction lives inside your post-trade process. Connectivity, workflow automation, real-time dashboards, and regulatory audit trails are all handled by the platform.



Built for T+1

Process confirmations with dozens of data points in seconds, supporting compressed settlement timelines T+1 requires.

How it works

Extract

The agent extracts specific data points. Each step is isolated, validated, and cross-checked for consistency.

Classify

The Agent identifies the confirmation type (commodity, FX, or other) and determines which fields to extract.

Receive

Xceptor monitors inboxes, SFTP location, or folders for incoming confirmations.

Validate

Once all extracted fields meet defined quality criteria, the Agent returns a response with confidence scores and document references for every field.

Deliver

Structured data is passed downstream in a standard JSON format, ready for matching, affirmation, or any other Xceptor workflow.

What changes when extraction is no longer a bottleneck

Extraction capacity that grows with trade volumes, not headcount

New counterparties onboarded in minutes, not days

Every extraction governed, traceable, and audit-ready

A post-trade operation that can clear in T+1

Ready to take confirmation processing to the next level?

[Learn more](#)

[Book a demo](#)

Xceptor