



Chioro — Feature Overview

by mgm technology partners | May 2026

Contents

1	Data Connectivity: Sources & Destinations	4
	Files & Formats	
	Databases & Platforms	
	Cloud & Remote Storage	
	API	
	Cross-Flow Data Reuse	
2	Data Processing Operations	6
	Visual Flow Builder	
	Transformation	
	Join, Union & Aggregation	
	Validation	
	Diff	
	External Service Calls	
3	AI Capabilities	8
	Agentic Flow Modelling	
	AI-Aided Attribute Mapping	
	AI Rule Generation from Natural Language	
	AI Connector Generation	
	AI Attribute Extraction	
	AI Chat Interface	
4	Rule Engine & Toolbox.....	10
	60+ Built-in Processing Functions	
	Composable Rules & Rule Blocks	
	JavaScript Extensibility	
	Real-Time Preview	
5	Data Exploration & Quality Assurance.....	11
	Data Analysis & Visualisation	
	Elasticsearch-Powered Data Exploration	
	Comparison View	
6	Automation, Operations & Enterprise.....	12
	Flexible Automation	
	Idempotent Execution	
	One-Time & Recurring Use	
	Git-Backed Flow Versioning	
	Multi-Flow Architecture	
	Role-Based Access Control	
	Notifications	
7	Deployment & Integration	14
	Cloud-Native SaaS	
	Extensible Platform	

Chioro – AI-augmented data processing

Chioro is an AI-augmented data cleansing and processing platform by mgm technology partners. It ingests data from any source and format, transforms and validates it through a combination of **deterministic** rule-based processing and targeted AI, and delivers clean, structured data to any target system – automatically. Chioro is not a PIM or MDM; it is a processing engine that sits between your sources and targets. It works equally well for recurring automated pipelines and one-time data transformations.

Chioro uses AI in two deliberate ways: as a **builder** that generates deterministic connectors, rule sets, and flow configurations from natural language or API documentation; and as a **processor** that handles per-record mapping when source data is too variable or ambiguous for static rules. In both cases, AI output is constrained by defined schemas and target structures.

1 Data Connectivity: Sources & Destinations

Connect to any data source and deliver to any target system — in any format, from anywhere.

	Sources	Destinations
Files & Formats	CSV, Excel, JSON, XML, BMEcat v1.2/2005, custom	CSV, Excel, JSON, XML, custom
Databases & Platforms	MySQL, PostgreSQL, SQL Server	PIMs, ERPs, CRMs, marketplaces, e-commerce platforms
Cloud & Remote Storage	FTP, SFTP, S3, Azure Blob, HTTP/HTTPS	FTP, SFTP, S3, Azure Blob, HTTP/HTTPS
API	Receive data via REST (inbound push or pull)	Push data via REST to any system (outbound)

Files & Formats

Ingest

Import CSV, Excel, JSON, XML, BMEcat v1.2/2005, and custom formats. No preconditions on data structure or content — import first, fix later.

Business benefit: *No more format rejection headaches. Whatever shape your data arrives in, Chioro accepts it.*

Deliver

Provide processed data as CSV, Excel, JSON, or XML. Field names, column order, and encoding are fully configurable to match the target system's requirements.

Business benefit: *Clean, correctly formatted files land exactly where downstream systems expect them, with no manual steps.*

Databases & Platforms

Ingest

Read directly from MySQL, PostgreSQL, and SQL Server via SQL queries. JDBC-based configuration, managed in the admin console.

Business benefit: *Pull data directly from ERP and CRM databases without manual exports.*

Deliver

Push processed data to PIMs, ERPs, CRMs, e-commerce platforms, and marketplaces via REST API or native integrations. For platforms that require deep integration — importing data

models, mapping attributes, and delivering in the platform's exact expected structure — purpose-built connectors are available; commercetools is a supported example.

Business benefit: *Chioro slots into the existing architecture without requiring changes in receiving systems and goes deeper where native integration is available.*

Cloud & Remote Storage

Ingest

Connect to FTP, SFTP, S3 (multiple providers), Azure Blob, and HTTP/HTTPS URLs. Wildcard file matching with change detection — only new or changed files trigger a run.

Business benefit: *Automated data pickup from any infrastructure, on-premise or cloud.*

Deliver

Write processed data to FTP, SFTP, S3, Azure Blob, or any HTTP/HTTPS endpoint. The same storage connections used for ingestion are available for delivery, with consistent configuration.

Business benefit: *Processed data lands in the right storage location automatically, without manual transfer steps.*

API

Ingest

Receive data via REST from any external system. Supports incremental APPEND and READY_TO_PROCESS signalling for streaming ingestion.

Business benefit: *Let upstream systems push data to Chioro as part of their own workflows.*

Deliver

Push processed records to any target system via REST API. Expose processed data for consumption by downstream systems on demand or as part of event-driven workflows.

Business benefit: *Processed data is immediately available to any system that needs it, on their schedule.*

Cross-Flow Data Reuse

Use the output of any operation in one flow as the input to another. Build modular, composable pipeline architectures without duplicating processing logic.

Business benefit: *Shared processing results flow between pipelines — no redundant logic, no data silos between flows.*

2 Data Processing Operations

Transform, cleanse, validate, and structure data with precision.

A schema defines the target structure — attributes, data types, mandatory fields, and allowed value lists — that processed data must conform to. Once defined or imported, it becomes the shared backbone for the pipeline: driving validation rules, auto-generating rule blocks, constraining AI extraction output, and guiding attribute mapping. A Product Type (e.g. "Electronic Device") is one example; multiple schemas can be organised into a full taxonomy covering any domain. Define the target structure once; everything downstream aligns to it automatically.

Visual Flow Builder

Drag-and-drop pipeline design. Connect operations with arrows; data flows along them. Test each step independently or run the full flow end-to-end.

Business benefit: *Non-developers can understand, validate, and modify data pipelines without code.*

Transformation

Rule-based attribute mapping, value transformation, text extraction, unit conversion, filtering, conditional logic, lookup table enrichment. 60+ built-in functions.

Business benefit: *Replace weeks of manual spreadsheet work with a configured, reusable rule set.*

Join, Union & Aggregation

Merge records from multiple sources by key, combine parallel data streams, or aggregate records into summary rows. Multiple flows or pipeline branches can feed into a single merge operation without requiring separate integration logic.

Business benefit: *Consolidate data from multiple suppliers, systems, or markets into a single clean output.*

Validation

Validate processed data against a schema — checking data types, mandatory fields, value lengths, and allowed value lists. Validation rules are derived directly from the schema definition, so there is no separate rule maintenance. Errors are surfaced explicitly, not silently skipped.

Business benefit: *Catch data quality problems before they reach downstream systems, with validation logic that stays in sync with the schema automatically.*

Diff

Compare two versions of a dataset and identify what has changed — additions, deletions, and modifications. Useful for detecting changes between supplier updates, verifying the effect of a transformation, or enabling incremental processing that only acts on new or changed records.

Business benefit: *Enable incremental processing and audit trails without storing full data copies.*

External Service Calls

Call any REST API inline within a flow — translation services (DeepL), media management (Cloudinary), AI/ML services, or internal systems. Results are returned directly into the pipeline and available as attributes for downstream operations.

Business benefit: *Enrich data automatically without building custom integration middleware.*

3 AI Capabilities

AI that builds reliable processing assets — and handles what rules cannot.

Across all AI operations, schemas serve as the shared guardrail: they define what the AI can produce, constrain its output to a known structure, and ensure that AI-generated results are always purposeful and auditable — never open-ended.

Agentic Flow Modelling

Provide a source data sample, a target structure, and a natural language description of the goal. Chioro's AI models the complete transformation pipeline. Review it, adjust it, then run it.

Business benefit: *Go from data sample to working pipeline in a fraction of the time of manual flow design.*

AI-Aided Attribute Mapping

Maps source data fields to target schema attributes at the structural level — resolving differences in column names, naming conventions, and field organisation between source and target. After a rule block is created based on a schema, the user triggers an Azure OpenAI API call to generate the mapping automatically. The result is human-reviewable and editable before it runs; the user applies their own judgement to the AI-proposed alignment.

Business benefit: *Onboard new data sources without manually mapping every field — the AI proposes the structural alignment, the user confirms.*

AI Rule Generation from Natural Language

Describe a transformation requirement in plain language: "extract the unit from the product title and convert to SI units". Chioro generates the rule set. It then executes deterministically on every record.

Business benefit: *Domain experts configure complex processing logic without needing to understand rule syntax.*

AI Connector Generation

Point Chioro at an API's endpoints and documentation. The AI agent reads the spec and generates a working, reusable connector. No developer required. The connector then runs deterministically.

Business benefit: *Connect to a new data source in hours, not weeks — without writing integration code.*

AI Attribute Extraction

Extracts structured attribute values from unstructured content — free-text descriptions, product titles, document fields — and populates target schema attributes with them. Each record is processed individually using OpenAI or Azure OpenAI, with output constrained by the

schema definition. Complements AI-Aided Attribute Mapping: where mapping resolves field alignment, extraction generates new structured data from content that has no direct source field.

Business benefit: *Surface structured, schema-conformant data from content that cannot be mapped with rules — turning free text into queryable, validated attributes.*

AI Chat Interface

Provides a conversational interface for exploring data, querying pipeline results, and configuring the flow in natural language. Queries run against the data currently in the flow at any processing stage.

Business benefit: *Business users can explore and interrogate data without SQL or technical tools.*

4 Rule Engine & Toolbox

Deterministic, extensible, and maintainable by business users.

Chioro supports two modes of rule authoring that can be freely combined: a no-code drag-and-drop configuration using built-in functions, and full JavaScript for cases where custom logic is needed. Both modes produce rules that execute deterministically and are versioned alongside the rest of the flow.

60+ Built-in Processing Functions

Text manipulation (case, trim, extract, regex, replace, split, join), number & unit handling (extract, convert units, format), list operations, lookup & decode tables, JSON/REST calls, structural operations.

Business benefit: *Cover the vast majority of cleansing requirements out of the box – no code needed.*

Composable Rules & Rule Blocks

Combine functions into rules; group rules into reusable blocks. Apply blocks across multiple transformations or import them from a schema.

Business benefit: *Define logic once, reuse everywhere. Consistency without repetition.*

JavaScript Extensibility

Write custom processing functions directly in the browser using JavaScript, or deploy via GitHub for team-managed tool libraries.

Business benefit: *Handle edge cases and business-specific logic without waiting for platform updates.*

Real-Time Preview

See the effect of any rule on a live sample of records instantly. Refine iteratively before committing to a full run.

Business benefit: *Build confidence in complex transformation logic before processing the full dataset.*

5 Data Exploration & Quality Assurance

Understand, analyse, and validate data at every stage of the pipeline.

Data Analysis & Visualisation

Automatic metadata detection, attribute-level overview and detail views, graphical quality dashboards, and category analysis.

Business benefit: Answer "what is the state of my data?" in seconds, not hours.

Elasticsearch-Powered Data Exploration

Data is indexed and searchable after every flow operation. Full-text search, attribute-level filtering, and statistical analysis — at any processing stage.

Business benefit: Diagnose data issues and understand source quality without writing queries.

Comparison View

Side-by-side comparison of input and output data at any operation, making the effect of transformations immediately visible.

Business benefit: Validate that transformations are doing exactly what was intended.

6 Automation, Operations & Enterprise

From manual one-time runs to fully automated enterprise data pipelines.

Flexible Automation

Run flows manually, on a time-based schedule (minute to yearly), via hotfolder (triggered when source data changes), or via REST API call. Conditional execution: only run if source data has changed since last execution.

Business benefit: Automate completely, or keep manual control — Chioro adapts to your process, not the other way around.

Idempotent Execution

The same input data with an unchanged flow always produces the same output, every time it is run. Deterministic and fully reproducible.

Business benefit: Enterprise-grade reliability — no surprises, no drift, no unexplained differences between runs.

One-Time & Recurring Use

Chioro is built for recurring automated pipelines and for one-time data transformations alike. Build the logic iteratively, refine until it is right, then execute with confidence. The flow remains available for re-use or archival.

Business benefit: Data migrations, consolidations, and clean-up projects benefit from the same tooling as daily operations.

Git-Backed Flow Versioning

Every flow configuration is stored in a Git repository. Create named backups at any point, restore to any prior state, and roll back changes instantly. Each agentic task represents a single Git commit.

Business benefit: Full audit trail and rollback for every pipeline configuration change.

Multi-Flow Architecture

Multiple independent flows coexist within a single Chioro client. Separate flows for separate data streams — different suppliers, markets, or processes.

Business benefit: Parallel data operations without interference or complexity.

Role-Based Access Control

Configure user permissions at the flow and operation level. Control who can view, edit, or execute each part of the pipeline.

Business benefit: *Enterprise-grade governance — the right people have the right access.*

Notifications

Configure notifications for flow execution events — success, failure, or data anomalies. Alerts can be directed to the right people or systems without requiring manual monitoring of the pipeline.

Business benefit: *Stay informed without manually monitoring runs.*

7 Deployment & Integration

SaaS-native, instantly productive, no installation required.

Cloud-Native SaaS

Chioro is hosted, managed, and updated by mgm technology partners. No installation, no infrastructure setup, no maintenance burden.

Business benefit: *Productive from day one. No IT project required to get started.*

Extensible Platform

Custom data formats, additional connectors, new flow operations, and custom tool functions can be added. Chioro is a platform, not a closed appliance.

Business benefit: *Handles today's standard requirements and tomorrow's edge cases.*