

PLATFORM WHITEPAPER



HATCH

Automated data collection, verification & decisioning

VERSION 3.0

A cloud-native, modular platform that turns static forms and manual review into real-time, rule-driven decisions for lending, leasing, and consumer finance — built end-to-end on a proprietary workflow and rule engine.

CONTENTS

Table of Contents

01	Abstract	03
02	Introduction	03
03	The Platform at a Glance	04
	<ul style="list-style-type: none">• Problem• Solution Pillars• Capabilities• Integrations	
04	How Hatch Works — End to End	06
05	Platform Architecture	07
	<ul style="list-style-type: none">• Administrator Portal• Point of Service• Orchestration Layer	
06	Core Engines & Capabilities	09
	<ul style="list-style-type: none">• Rule Engine• Open Banking• Fraud• KYC• Workflow Builder• Audit Ledger• Payments• Servicing• Analytics & API	
07	Technology Stack	13
08	Scalability & Performance	14
09	Security & Compliance	14
10	Deployment Models	15
11	Use Cases	15
12	Conclusion	18

SECTION 01

Abstract

In lending, leasing, and consumer finance, the old way of collecting data — static forms, manual entry, and serial back-office review — is slow, error-prone, and increasingly exposed to fraud. **Hatch replaces that model with a single, automated platform.**

Hatch is a modular, cloud-native web application that lets a business stand up a complete application-to-decision workflow in minutes. A drag-and-drop builder removes the technical barrier to creating complex logic; a proprietary rule engine evaluates each applicant against configurable criteria in under a second; and live integrations to open banking, credit bureaus, identity-verification, payment, and document services let the platform gather and act on real data automatically. Every decision is written to an immutable, cryptographically-hashed audit ledger.

The result is faster approvals, lower default and fraud exposure, demonstrable compliance, and dramatically reduced operational cost. Hatch also extends these capabilities directly to consumers through a B2C product that turns previously inaccessible verification and analysis into a single click. This whitepaper details the platform's architecture, its core decisioning engines, its technology stack, and the industry use cases it transforms.

SECTION 02

Introduction

Industries that run on applicant data — consumer and business lending, mortgage servicing, debt recovery, auto and residential leasing — share a common problem. Their intake processes are built on rigid forms and manual keying, the workflows behind them are hard to integrate with existing websites and data sources, and the people who depend on the outcome (both staff and applicants) wait far longer than they should for a decision. Consumers, meanwhile, rarely have direct access to the verification and analysis tools that determine those decisions.

Hatch addresses all three problems at once. Businesses configure and deploy workflows through an intuitive drag-and-drop interface, connect them to an existing site with a few clicks, and let the platform's automation handle data collection, verification, decisioning, and downstream actions such as contract delivery and disbursement. For consumers, Hatch's B2C product generates a shareable, personalized application link that returns a real-time data aggregation screen — including instant bank-verification (IBV) analysis — democratizing access to information that used to require an intermediary.

SECTION 03

The Platform at a Glance

3.1 The Problem

Traditional data collection introduces bottlenecks at every step. Manual entry creates errors and delays; rigid forms cannot adapt to a specific product or risk policy; and connecting a workflow to an existing website or data source is a technical project in its own right. The cost of legacy review — staff, third-party data pulls, legal and regulatory overhead — keeps rising, while subprime and high-volume segments make accurate, real-time risk assessment harder than ever. Consumers are left without access to the verification data that decisions are based on.

3.2 What Hatch Does — Four Pillars



No-code workflow setup

A drag-and-drop builder lets anyone create, modify, and deploy complex workflows in minutes — no engineering cycle required.



Advanced automation

A proprietary rule engine drives inter- and intra-system integrations, automating decisions and data flow with consistency and accuracy.



Simple connectivity

Workflows link directly to an existing website with a few clicks, enabling seamless data collection and real-time updates.



B2C data access

Consumers generate a unique link and instantly receive a data-aggregation screen with IBV analysis — verification in one click.

3.3 Product Capabilities

- ◆ **Modular design** — adaptable to lending, leasing, or any data-driven vertical; configure once per product line.
- ◆ **Data integration** — IBV, KYC, credit bureaus, and open banking feed accurate, real-time information into every workflow.
- ◆ **Custom workflows** — design intake and decision logic that mirrors your operational process exactly.

SECTION 03.4

List of Integrations

Hatch ships with dedicated connectors that handle authentication, provider-specific rate limits, error retries, and schema normalization — so varied external responses are translated into one unified internal format before they reach the rule engine.

CATEGORY	SERVICES	ROLE IN THE WORKFLOW
Credit bureaus	TransUnion, Equifax, Experian	Real-time credit scoring and financial-history retrieval.
Open banking / IBV	Flinks, Plaid	Bank-account linking and transaction-data retrieval.
Payments	Plaid, Vopay, Zoomrails, EFT Canada	Push payments, wire, and EFT disbursement on approval.
Identity / KYC	OnDato	Document authenticity and facial-liveness verification.
Documents & e-sign	PandaDoc, DocuSign	Contract generation and signature lifecycle.
CRM	Pipedrive	Webhook-driven deal-stage synchronization.

EXTENSIBLE BY DESIGN

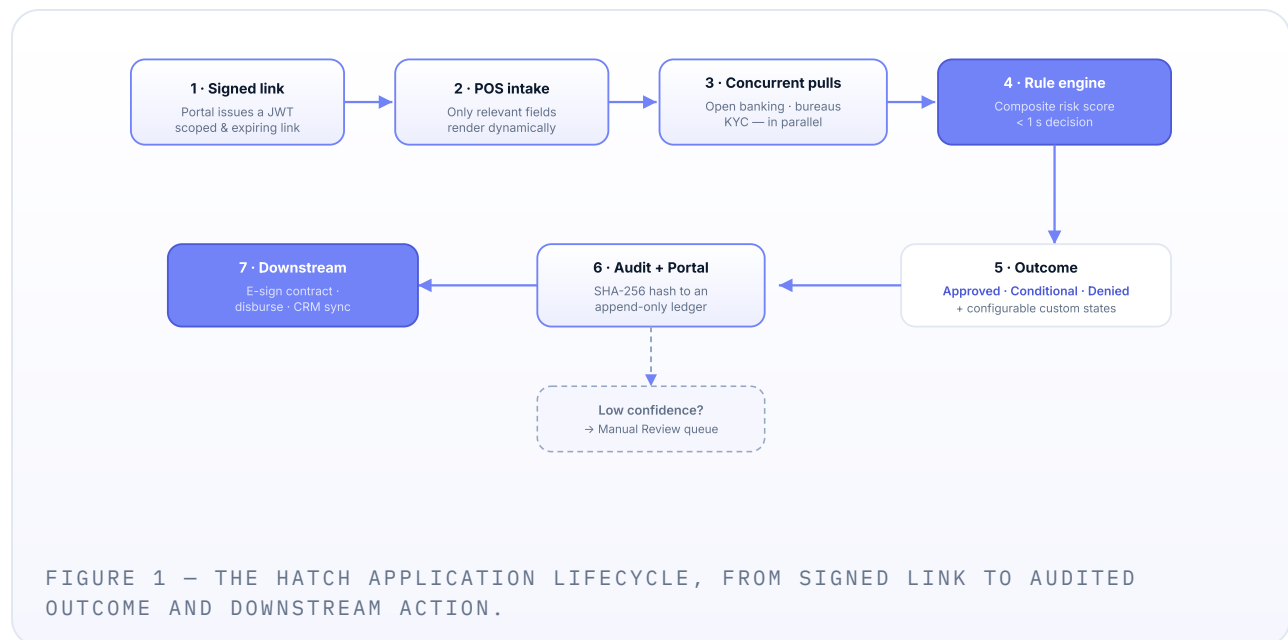
The connector framework is open: new bureaus, banking providers, court databases, land and personal-property registrars, valuation services, and background-check providers can be added as standardized adapters without changing core decision logic.

Hatch exposes two surfaces to the outside world: the **Administrator Portal**, where lenders and operators evaluate, approve, and monitor applications, and the **Point of Service (POS)** app, which carries the applicant from first input through final approval. The sections that follow trace a single application across both.

SECTION 04

How Hatch Works — End to End

A single Hatch application moves from a signed link to a fully-audited decision — and, where configured, to a signed contract and disbursed funds — without manual hand-offs.



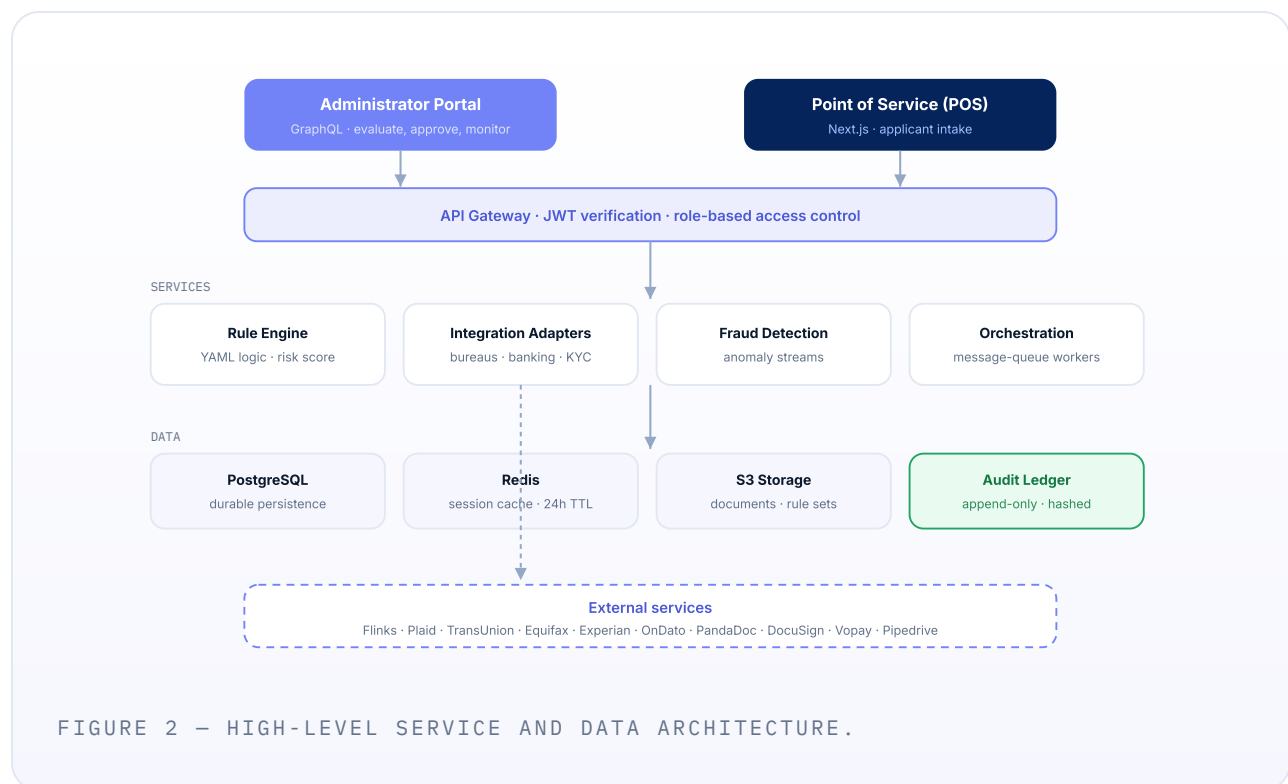
The seven steps

- ◆ **1. Signed link.** The Portal embeds workflow metadata — product type, language, expiry — into a JSON Web Token and issues a digitally-signed, scoped application link.
- ◆ **2. POS intake.** When the applicant opens the link, the Next.js POS renders only the fields relevant to that product and posts responses to the backend.
- ◆ **3. Concurrent retrieval.** In parallel, micro-services call open-banking providers, credit bureaus, and KYC services to fetch live financial and identity data.
- ◆ **4. Evaluation.** The rule engine scores all inputs against the workflow's configured logic in under a second (P95 < 800 ms).
- ◆ **5. Outcome.** The decision — Approved, Conditional, Denied, or a custom state — is returned to the Portal.
- ◆ **6. Audit.** A SHA-256 hash of the inputs, fetched data, and rule-set version is written, with a UTC timestamp, to an append-only ledger.
- ◆ **7. Downstream.** On approval, Hatch can generate and send a contract for e-signature, disburse funds, and update the CRM — automatically.

SECTION 05

Platform Architecture

Hatch is built as a set of independently-deployable micro-services behind a unified API. Two front-ends — the Administrator Portal and the Point of Service app — sit on top of a shared services layer: the rule engine, integration adapters, fraud detection, and the orchestration workers that sequence each workflow. A relational store, an in-memory cache, object storage, and an append-only ledger form the data layer.



5.1 Administrator Portal

The Portal is the operator's central hub — the place lenders, lessors, and processors evaluate, approve, and monitor every application. It houses the analytics framework, the rule engine, the credit-data integration layer, and fraud detection, and surfaces a "Manual Review" queue whenever a decision's confidence score falls below a configured threshold. It is built on a modular architecture designed for scalability, security, and seamless integration with external data sources.

5.2 Point of Service (POS) App

The POS is the applicant's experience. Built on Next.js for a high-performance, low-friction frontend, it verifies the authenticity of each application link before rendering, then dynamically

5.3 Orchestration & Automation Layer

Multi-stage workflows are sequenced by Hatch's own orchestration layer. Workflow definitions authored in the visual builder are serialized to a versioned JSON/YAML structure and executed directly by the platform's proprietary engine — there is no dependency on any external process-orchestration product. Message-queue workers drive each step: data fetches, decision forks, manual tasks, and downstream actions are dispatched, retried, and tracked as discrete jobs, which is what allows individual stages to scale independently under load.

VALIDATION, COMPLIANCE & AUDIT TRAILS

The automation layer also runs application validation and compliance checks at each stage, and emits the events that populate the immutable audit ledger — minimizing manual intervention while keeping every step traceable.

Real-time monitoring & logging

A monitoring service provides real-time logging and alerting on system performance, so issues are detected and addressed promptly and the platform's behavior is observable end-to-end.

Role-based access control

A comprehensive RBAC system governs permissions across the Portal, ensuring only authorized personnel reach sensitive data and administrative functions.

RBAC

GRANULAR PERMISSIONING

24_h

RESUMABLE SESSION TTL

Real-time

MONITORING & ALERTS

Because workflows pause and resume against the Redis session cache, an applicant can leave a partly-completed application and return to it later without data loss — a small detail that materially lifts completion rates.

SECTION 06

Core Engines & Capabilities

6.1 Configurable Rule Engine & Real-Time Decisioning

The rule engine is the heart of the platform. It evaluates each applicant's combined data against a configurable rule set and returns a decision in under a second. Each application link carries a product identifier — for example MORTGAGE or AUTO-LEASE — which the engine resolves to a corresponding rule set, so mortgage criteria (debt ratios, valuations) never mix with auto-lease logic (residual calculations).

Decisions are driven by a **composite risk score**: a logistic-regression model trained on historical default data applies weighted coefficients across signals such as credit score, debt-to-income ratio, and income stability. Those weights — for instance 40% credit, 30% DTI, 30% income stability — live in external definition files, so business teams adjust the model instantly, without a code deployment. Rule sets are **hot-reloaded**: replacing a definition takes effect on new applications immediately, with no restart or redeploy, which turns a compliance fix or policy change from a release into a configuration edit.

Decision latency	P95 < 800 ms
Outcomes	Approved · Conditional · Denied · custom
Model definition	External, version-controlled, hot-reloaded
Manual-review trigger	Confidence below configurable threshold

6.2 Open Banking & Financial-Data Aggregation

Hatch's open-banking adapter ingests and normalizes hundreds of disparate transaction formats into ISO 20022-compliant records — capturing over 2,300 distinct data points for personal accounts and more than 1,800 for business accounts. Distributed routines compute rolling averages over 30/60/90/120/365-day windows, identify recurring salary or revenue deposits, tally non-sufficient-fund events, and calculate a dynamic cash-flow stability score. These insights feed the rule engine directly, letting risk models evaluate real income and expense behavior rather than relying on a static credit score alone.

6.3 Fraud Detection

A dedicated fraud module subscribes to event streams and applies device fingerprinting, IP/ASN lookups, velocity checks, and AI-driven analysis of uploaded identity documents. When a heuristic fires — a mismatched geolocation, rapid repeat submissions — the system tags the case

6.4 Identity Verification (KYC)

Hatch integrates third-party KYC services to perform document-authenticity checks and facial-liveness verification. The pass/fail result becomes a variable in the rule engine, enabling automated declines on a failed check or routing to manual review when results are inconclusive.

6.5 No-Code Workflow Builder

A browser-based editor lets business users assemble workflows visually — dragging and connecting nodes such as data fetch, decision fork, and manual task — and edit decision parameters in place. The editor serializes the resulting graph and parameters to JSON/YAML, which the orchestration engine consumes directly. The effect is that routine business changes no longer require developer cycles.

6.6 Immutable Audit Ledger

On completing a decision, Hatch concatenates the JSON payload — applicant data, fetched data, and rule-set version — and computes a SHA-256 hash. That hash, with a precise UTC timestamp, is written to a write-once ledger table that cannot be updated or deleted. The SHA-256 digest of each rule-set version is stored alongside the decision, so auditors can always match an outcome to the exact logic that produced it. The ledger supports SOC 2 and regulatory audits and makes every decision provably traceable.

MANUAL INTERVENTION IS AUDITABLE TOO

When an underwriter overrides or augments an automated result, the rationale and user ID are recorded — producing an explicit trail of every human touch alongside the automated record.

6.7 Cost-Aware Orchestration

Not every applicant needs every data pull. When the composite risk score indicates clear high risk, the orchestration layer suppresses non-essential bureau calls — saving roughly \$3 per run in fixed data costs, on average, without compromising the integrity or speed of the core decision. At volume, that selective suppression is a direct, recurring margin gain.

6.8 Disbursement & Payments

On an approved decision, Hatch can trigger a payment instruction in the same workflow, supporting any of three rails:

- ◆ **Push payment** to a recipient's phone number or email, using pre-linked credentials — no account details required.
- ◆ **Wire transfer** formatted and submitted over banking rails.

6.9 Loan Servicing & Intelligent Rescheduling

Hatch's data-integration layer connects to its own or external loan-servicing platforms to retrieve an applicant's full installment schedule — each due date, principal and interest amount, and outstanding balance — normalizing it into the same ISO 20022-style records as transactional data. The platform continuously monitors for missed payments and, when an expected debit fails to clear, triggers an automated rescheduling workflow.

Using the borrower's historical cash-flow patterns, the rule engine selects new payment dates aligned with expected fund availability to maximize successful collection. Rescheduled plans can be set to any interval — daily, weekly, biweekly, bimonthly, or monthly — as well as interest-only, principal-only, or a deferred grace period. The updated plan feeds back into the workflow, driving reminders, refreshed amortization tables, and downstream risk rules.

6.10 Documents & E-Signature

On approval, a document service compiles a PandaDoc or DocuSign template, merges in applicant details, and sends the signature request. Signed-status webhooks update the Portal and store the fully-executed document in secure file storage — automating the entire signature lifecycle.

6.11 CRM & Webhooks

At each workflow stage — Submitted, Under Review, Approved — Hatch emits a JSON webhook to a connected CRM such as Pipedrive, advancing the deal stage automatically. Sales and operations stay aligned without manual data entry and without the sync errors it produces.

6.12 Analytics & Reporting

Hatch aggregates event logs and presents KPIs — application counts, approval rates, mean decision latency — through interactive dashboards that render in the browser without page reloads. Users filter by date range, product type, or region and export to CSV for deeper analysis.

6.13 Open API

The rule-engine service exposes a REST API so external systems can initiate and monitor workflows programmatically. A client posts a workflow key and payload, receives a unique run identifier and a streaming endpoint, and can track progress and retrieve the final outcome in real time.

Workflow trigger	REST · workflow key + payload
Progress tracking	Run ID + real-time stream
Dashboard refresh	In-browser, no page reload

SECTION 06 · SUMMARY

Capabilities, at a Glance

CAPABILITY	WHAT IT DELIVERS
Rule engine	Product-aware, sub-second decisions from externally-configured, hot-reloadable logic.
Open banking	2,300+ / 1,800+ normalized data points; cash-flow stability scoring.
Fraud detection	Device, network, velocity, and document-image anomaly detection.
KYC	Document authenticity + liveness as a decision variable.
Workflow builder	Visual, no-code authoring serialized to executable JSON/YAML.
Audit ledger	Write-once SHA-256 records tied to rule-set version; SOC 2-ready.
Cost control	Risk-aware suppression of unneeded bureau pulls (~\$3/run saved).
Payments	Push, wire, and EFT disbursement with sub-5-second settlement.
Servicing	Missed-payment detection and cash-flow-aligned rescheduling.
Documents	Automated contract generation and e-signature lifecycle.
CRM	Stage-by-stage webhook synchronization.
Analytics & API	Interactive dashboards and a REST/streaming workflow API.

ONE PLATFORM, CONFIGURED PER PRODUCT

Every capability above is governed by the same workflow and rule-engine model. A new product line is a new configuration — not a new build.

SECTION 07

Technology Stack

Hatch is engineered for performance, maintainability, and horizontal scale. Responsibilities are divided cleanly across services: the Next.js POS handles applicant input, NestJS services call external APIs and host the rule engine, message-queue workers orchestrate multi-stage tasks, and a GraphQL layer serves the Portal. Persistence and caching are split between a relational database and an in-memory store.

LAYER	TECHNOLOGY	PURPOSE
Frontend	Next.js	High-performance Portal and POS interfaces across platforms.
Backend	NestJS	Scalable, maintainable services for complex business logic and integrations.
Orchestration	Message-queue workers	Sequence and retry multi-stage workflow jobs; proprietary engine — no external BPM dependency.
Relational data	PostgreSQL	Durable application, decision, and audit-ledger persistence.
In-memory data	Redis	Session caching and real-time state (24-hour TTL) for resumable workflows.
Storage & hosting	AWS · S3	Cloud hosting and scalable, secure object storage for documents and rule sets.

ARCHITECTURE NOTE

Earlier iterations of Hatch evaluated an external process-orchestration engine. The platform now runs its own orchestration layer end-to-end, removing that dependency and consolidating workflow definition, execution, and versioning inside Hatch.

SECTION 08

Scalability & Performance

Hatch follows a service-oriented architecture: each service is a containerized image that scales horizontally and independently, with autoscaling driven by real-time workload metrics — CPU load and job-queue depth — to hold sub-second response times during traffic surges and scale down to cut cost when idle. Rigorous end-to-end testing with zone redundancy ensures reliable, consistent performance under high transaction volumes.

Auto

DEMAND-BASED SCALING

Zone

REDUNDANT E2E TESTING

SOA

INDEPENDENT SERVICE SCALING

SECTION 09

Security & Compliance

Data protection

All network traffic uses TLS 1.3 with strong ciphers. Personally-identifiable and financial fields are encrypted at rest with AES-256 under AWS KMS-managed keys, and uploaded files such as bank statements are stored with server-side encryption and strict access controls.

Link & access security

Application links embed a digital signature over a JWT payload, with an expiry claim (default 24 hours) and a scope claim limiting which workflow stages the link can reach; expired, tampered, or replayed links fail verification and are rejected. Inside the Portal, role-based access control governs every sensitive function.

Regulatory compliance

Continuous monitoring and auditing maintain compliance with GDPR, CCPA, and PIPEDA. On a data-erasure request, the platform pseudonymizes or deletes personal data while preserving anonymized analytical metrics — satisfying privacy obligations without losing business intelligence.

In transit	TLS 1.3
At rest	AES-256 · AWS KMS
Audit	Write-once SHA-256 ledger
Frameworks	GDPR · CCPA · PIPEDA · SOC 2

SECTION 10

Deployment Models

Hatch is delivered as infrastructure-as-code — container images for the POS, rule engine, and Portal services, plus deployment and provisioning scripts — so the full environment reproduces consistently on any major cloud. For organizations with strict regulatory or network-isolation requirements, a pre-configured machine image bundles the database with its schema migrations, the in-memory store, the orchestration engine, the backend services, and the frontend, enabling on-premises deployment.

Cloud SaaS

Fully-managed, autoscaling deployment on AWS — the default for most operators.

On-premises / isolated

Self-contained image for regulated or air-gapped environments, deployable in the customer's own infrastructure.

SECTION 11

Use Cases

The same platform, configured per product line, transforms six core workflows.

11.1 Transforming Consumer Lending

Hatch automates the loan application process end-to-end — instantly gathering and analyzing financial data, returning real-time credit-risk assessments, and predicting default risk — while integrating with existing loan-processing systems.

IMPACT

Faster approvals, lower default rates, higher operational efficiency.

KEY FEATURES

Automated data collection · real-time decisioning · ML risk management · seamless integration.

11.2 Optimizing Mortgage Servicing

Hatch streamlines mortgage servicing for servicers and solicitors — automating document processing, integrating real-time data on borrower standing and property values, and monitoring compliance — with workflows customized to each servicer.

11.4 Empowering Business Subprime Lending

For lenders facing high default rates and thin applicant data, Hatch builds ML-powered credit-scoring models tuned to subprime business risk, automates underwriting, and monitors loan performance in real time with default alerts.

IMPACT

Improved default rates, greater lending efficiency, enhanced profitability.

KEY FEATURES

Advanced credit scoring · automated underwriting · risk monitoring & alerts · customizable criteria.

11.5 Enhancing Auto Leasing & Financing

Hatch removes the manual entry and slow credit checks that bog down auto leasing — instantly evaluating creditworthiness, providing real-time risk assessments, and automating the application process from document collection through approval.

IMPACT

Faster approvals, reduced defaults, higher throughput.

KEY FEATURES

Automated credit assessment · risk-management tools · streamlined application · customizable financing plans.

11.6 Transforming Residential Leasing

Hatch enhances every step of residential leasing — automating tenant screening with comprehensive data (credit, rental history, risk), streamlining document, communication, and rent-collection workflows, and providing real-time vacancy management.

IMPACT

Faster tenant placement, better tenant quality, reduced vacancies.

KEY FEATURES

Automated screening · streamlined leasing workflows · real-time vacancy management · customizable lease agreements.

A CONSISTENT PATTERN

Across all six, the mechanics are the same — collect data through a signed link, enrich it with live integrations, decide in under a second against configurable logic, and act automatically — differing only in the rule set and the data sources each product requires.

SECTION 12

Conclusion

Hatch represents a transformative leap in how businesses and consumers manage data in lending, leasing, and beyond.

By replacing static forms and manual review with a single automated platform, Hatch collapses the distance between an applicant's first input and a fully-audited, acted-upon decision. Its drag-and-drop builder removes the technical barrier to complex workflows; its proprietary rule engine decides in under a second against logic that business teams control directly; its live integrations gather and act on real financial, identity, and behavioral data; and its immutable ledger makes every decision provably traceable.

The versatility of the platform shows across its surfaces — from the Administrator Portal that centralizes and optimizes the operator's workflow, to the Point of Service app that makes the applicant's journey effortless, to the B2C product that puts verification data directly in consumers' hands. Built on a service-oriented architecture with security, compliance, and ease of use at its core, Hatch scales with demand and deploys equally well in the cloud or on-premises.

As industries demand more from their data-management solutions, Hatch is positioned to lead — delivering speed, accuracy, lower cost, and demonstrable compliance to businesses, and unprecedented access to their own data to consumers. It is a comprehensive solution built for the present and the future of data-driven finance.



HATCH LABS · HATCHLABS.TECH · WHITEPAPER V3.0

This document is confidential and intended for the recipient only. Figures and capabilities described reflect the Hatch platform as designed and may evolve.