

AUTHENTIC

Authentic AI. Augmented Outcomes.

Agentic AI for the Modern Enterprise

Fundamentals · Use Cases · Readiness · Roadmap

Prepared for
Prospective Client

HOF T &
HUB OF FINANCE TRANSFORMATION

From Automation to Autonomy

Three eras of work automation

Robotic Process Automation

Followed scripted rules. Broke when forms changed.

Generative AI Assistants

Drafted answers. A human still ran the workflow.

Agentic AI

Plans, decides and executes complete workflows under guardrails — measured by outcomes. not clicks.

70%+

of finance and back-office tasks are agent-addressable

30–60%

cycle-time reduction in early agentic deployments

24×7

execution capacity, scaling without linear headcount growth



The question is no longer “can AI help?” — it is “which work should we delegate first, and how do we govern it?”

What we will cover

1	Foundations	What an agent is, how it works, and how it differs from RPA and chatbots
2	Where to deploy	Identifying agent-fit processes and strong starting points
3	Use cases & ROI	Eleven business domains with measurable outcomes
4	Readiness	Data, process, systems, organisation and governance prerequisites
5	Responsible deployment	Human oversight, trust, monitoring and the 90-day path

What is an autonomous AI agent?



An autonomous AI agent is a software entity that perceives a goal, reasons about how to achieve it, plans a sequence of steps, executes them by calling tools and systems, and reflects on the outcome — all while operating within defined guardrails.



Goal-oriented

Pursues an outcome, not a script. Re-plans when reality changes.



Tool-using

Calls APIs, ERPs, browsers and databases — the same tools your team uses.



Self-correcting

Detects errors, retries, escalates to a human when uncertain.



Autonomous within bounds

Operates inside policy, approval and risk-tolerance guardrails.

Where agents differ from what came before

Capability	Traditional RPA	Generative Assistant	AI Agent
Handles unstructured input	No — breaks on layout change	Yes — drafts and summarises	Yes — and acts on it
Plans multi-step workflows	Pre-coded steps only	Suggests steps to user	Plans and re-plans dynamically
Executes across systems	Yes (UI scripts, brittle)	No — outputs text only	Yes via tools and APIs
Adapts when context changes	No	Limited	Yes — reasons over new state
Recovers from errors	Halts; needs developer	Apologises; user retries	Detects, retries, escalates
Improves over time	Manual rule updates	Better prompts	Learns from feedback and outcomes

How agents plan and execute end-to-end



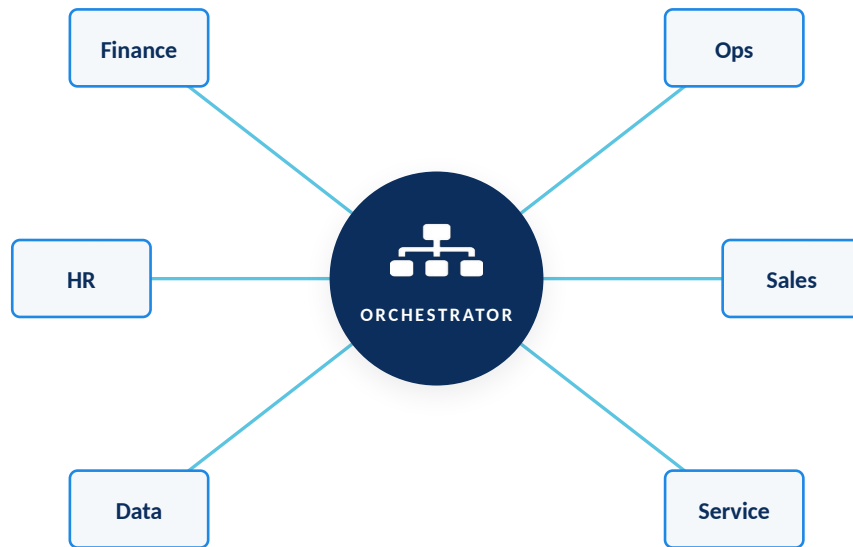
Single agents become teams

From soloists to specialists

Real workflows span functions. A multi-agent system divides labour across specialist agents, coordinated by an orchestrator that owns the goal, sequences work, resolves conflicts and reports outcomes.

Why teams of agents

- Each agent stays focused on a narrow domain it does well
- Tasks run in parallel, not in a single brittle script
- Failures stay local; the orchestrator can re-route
- Governance and audit are layered per agent and per action



A goal flows in. The orchestrator dispatches work to specialists, gathers results, and returns one coherent outcome.

What makes a process agent-fit



High volume, repeatable

Enough monthly transactions to amortise design and oversight cost



Multi-system and multi-format

Spans ERPs, email, PDFs, spreadsheets — RPA's natural blind spots



Judgement-light decisions

Rules can be expressed clearly; edge cases are bounded and known



Measurable outcomes

Cycle time, accuracy, exception rate, days saved — quantifiable



Tolerable failure mode

Errors are recoverable, reversible, or caught by a downstream check



Stable but not trivial

Worth automating: not so simple that a script suffices, not so creative that humans must own it

Identifying strong starting points



Score each candidate (1-5)

1. Volume of transactions
2. Cycle-time pain today
3. Data availability & quality
4. Tool / API accessibility
5. Tolerance for occasional error
6. Owner motivated to change

Finance — FP&A, Accounting, Treasury



FP&A

Planning, forecasting & narrative

Agentic use cases

- Continuous rolling forecast updates from drivers
- Variance-to-budget commentary auto-drafted
- Scenario simulation across price, volume, FX
- Board-pack narrative generation

OUTCOME / ROI

Forecast cycle from 10 days → 2 days; up to 60% less time on commentary; faster CFO insight loop



Accounting

Close, controls & compliance

Agentic use cases

- Invoice extraction, coding & 3-way match
- Bank and intercompany reconciliations
- Journal posting with policy-aware controls
- Close-task tracking & exception triage

OUTCOME / ROI

Close compressed by 30–50%; ~70% of routine recs auto-cleared; control evidence captured automatically



Treasury

Cash, FX & risk

Agentic use cases

- Daily cash position from multi-bank feeds
- 13-week cash forecast with driver-level traceability
- FX exposure tracking and hedge proposals
- Payment run validation & fraud screening

OUTCOME / ROI

Cash visibility within hours, not days; higher hedge effectiveness; fraud risk caught before payment release

People & front office — HR, Sales, Marketing



HR

Talent, ops & employee experience

Agentic use cases

- Job-spec drafting and screening shortlists
- Onboarding workflow & document handling
- Policy & benefits Q&A for employees
- Attrition signals from engagement data

OUTCOME / ROI

Time-to-shortlist down 40–60%; HR ticket deflection 50%+; faster, more consistent onboarding



Sales

Pipeline, prep & follow-through

Agentic use cases

- Lead enrichment, scoring & next-best action
- Meeting prep briefs from CRM, news, filings
- Proposal & RFP first-draft generation
- Pipeline hygiene and forecast accuracy

OUTCOME / ROI

Seller admin time cut 30–40%; meeting-to-opportunity conversion uplift; cleaner forecast



Marketing

Demand, content & analytics

Agentic use cases

- Campaign briefs and multi-channel content drafts
- Audience segmentation & journey orchestration
- Performance attribution across channels
- Competitor & market scanning

OUTCOME / ROI

Content production capacity 2–3× per FTE; campaign launch time halved; sharper attribution decisions

Supply chain & operations



Procurement & Inventory

Source-to-pay & stock optimisation

Agentic use cases

- Supplier discovery, scoring & risk monitoring
- RFP drafting and bid evaluation support
- PO matching, exception clearing, contract checks
- Demand-aligned reorder and safety-stock tuning

OUTCOME / ROI

PO cycle time down 40%+; working capital release 5–15%; fewer stock-outs and write-offs



Logistics

Plan, route & track

Agentic use cases

- Carrier selection and rate negotiation prep
- Shipment routing & ETA monitoring
- Exception handling for delays and damages
- Customs and documentation drafting

OUTCOME / ROI

On-time delivery uplift 5–10pp; freight cost optimised 3–8%; fewer manual track-and-trace touches



Operations & Customer Service

Run the business, serve the customer

Agentic use cases

- Production scheduling and OEE diagnostics
- Quality issue triage from logs and reports
- Tier-1 customer queries with system actions
- Complaint resolution with policy-aware decisions

OUTCOME / ROI

Operator productivity +15–25%; first-contact resolution +20–40%; sustained CSAT improvement

Knowledge & decisions



Document Management & Admin

Read, classify, act, file

Agentic use cases

- Contract review with clause and risk extraction
- Inbox triage, routing & response drafting
- Policy and SOP search across repositories
- Meeting minutes, action capture and follow-up

OUTCOME / ROI

Document handling time down 50-70%; faster contract turnaround; institutional knowledge becomes searchable



Advanced Analytics

Decision support, not just dashboards

Agentic use cases

- Natural-language Q&A across financial and ops data
- Anomaly detection with root-cause hypotheses
- Predictive models for demand, churn and credit
- Always-on KPI monitoring with proactive alerts

OUTCOME / ROI

Insight-to-action cycle from days to minutes; model-driven decisions in pricing, credit, demand; analyst capacity reallocated to strategic work

How to evaluate agent performance

EFFICIENCY

- Cycle time per case
- Cost per transaction
- Throughput vs baseline
- Hours redeployed to higher value

QUALITY & ACCURACY

- Error / rework rate
- First-time-right %
- Exception escalation rate
- Audit and policy compliance

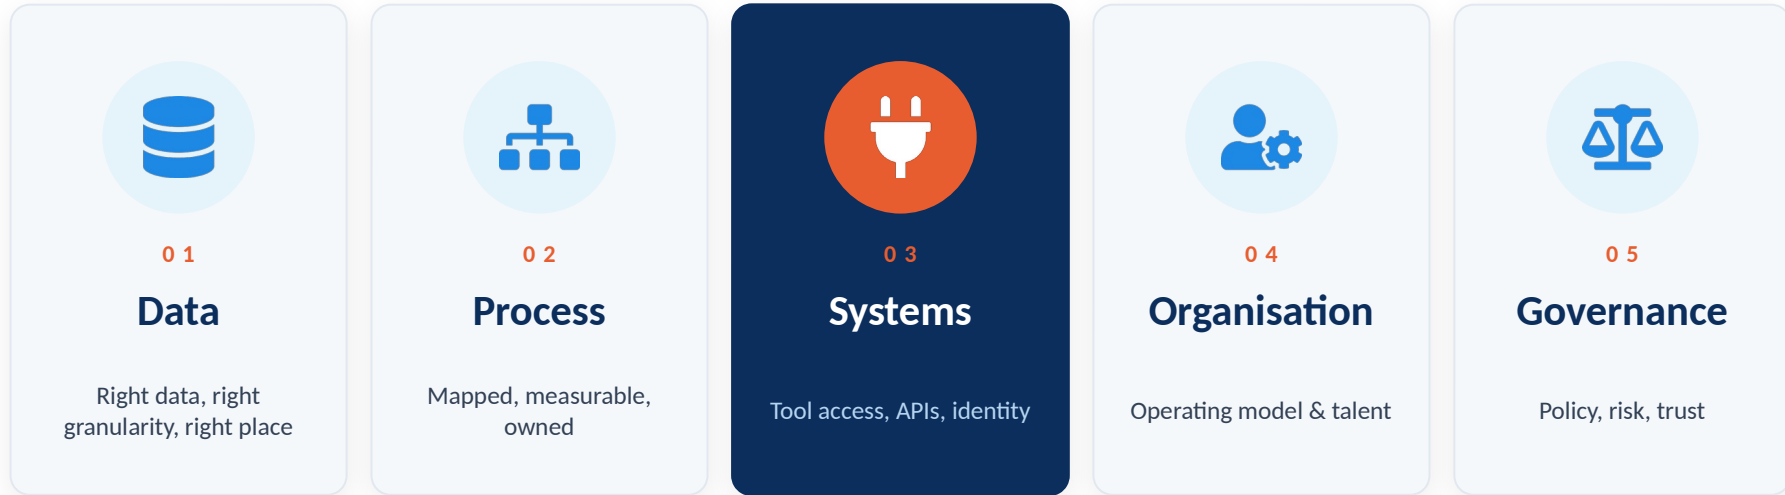
STRATEGIC OUTCOMES

- Revenue or margin uplift
- Working capital released
- Customer NPS / CSAT change
- Risk events prevented

ADOPTION & TRUST

- % of cases run end-to-end by agent
- Human override rate (and trend)
- User satisfaction with agent output
- Time to onboard new use cases

Five pillars for agentic success



Each pillar is unpacked in the next four slides.

Data readiness



PILLAR 1

Agents only act as well as the data they can see. Treat data as the foundation, not an afterthought.

FOUNDATIONS

- Single source of truth for finance and operations
- Lakehouse / warehouse with curated semantic layer
- Master data: customers, vendors, products, accounts
- Document repositories indexed and searchable

QUALITY & ACCESS

- Defined data owners and quality SLAs
- Lineage and freshness visible to consumers
- Granular role-based access controls
- PII and sensitive data classified and masked

AI-READINESS

- Vector indexes over policies and unstructured data
- Real-time signals where decisions are time-critical
- Feedback loops capturing agent outcomes
- Synthetic data for safe testing of edge cases

Process & systems readiness



PILLARS 2 & 3

Agents need processes worth automating and systems they can actually reach.

PROCESS DESIGN

- End-to-end maps with inputs, outputs, owners
- Decision rules made explicit and testable
- Exceptions catalogued with handling logic
- Baseline KPIs captured before deployment

SYSTEMS & INTEGRATION

- API access to ERP, CRM, banking, HRIS, ticketing
- Identity and service accounts for agent actions
- Sandbox environments mirroring production
- Logging across every agent-system interaction

ARCHITECTURE

- Orchestration layer for multi-agent coordination
- Tool registry with authorisation per action
- Observability: traces, prompts, decisions, results
- Cost and rate-limit controls by use case

Organisation & talent readiness



PILLAR 4

Agentic AI changes the operating model. Plan the people side as deliberately as the technology.

OPERATING MODEL

- AI Centre of Excellence with clear remit
- Federated business owners for each use case
- Joint product teams: business + data + AI
- Funding model that survives quarterly cycles

ROLES & SKILLS

- Process owners trained to design with agents
- AI engineers, prompt designers, eval specialists
- Risk, compliance and audit upskilled on AI
- Front-line teams trained as agent supervisors

WAYS OF WORKING

- Decisions on what to delegate vs retain in-house
- Career paths that survive automation of tasks
- Communication plan to manage uncertainty
- Recognition for redeployment, not just headcount cuts

Governance, risk & ethics



PILLAR 5

Without governance, agentic AI scales mistakes as fast as it scales work. With it, trust compounds.

POLICY & STANDARDS

- AI usage policy aligned to local regulation
- Approval gates by risk tier of use case
- Acceptable-use rules for data and models
- Vendor and model selection standards

RISK CONTROLS

- Pre-deployment evaluation: accuracy, bias, security
- Action-level authorisation limits
- Logging, audit trails and forensic replay
- Incident playbook with rollback procedures

ETHICS & ACCOUNTABILITY

- Named accountable executive per use case
- Transparency to users when an agent is involved
- Fairness checks for decisions affecting people
- Independent review for high-impact deployments

Human oversight & trust building

Oversight scales with risk

1

Human-in-the-loop

Every agent action requires human approval

2

Human-on-the-loop

Agent acts autonomously; humans review samples & exceptions

3

Human-out-of-the-loop

Fully autonomous within bounds; humans review aggregates

Building and sustaining trust



Transparency

Show what the agent did and why



Verifiable outputs

Citations, calculations, source links



Continuous monitoring

Drift, accuracy and bias dashboards



Feedback channels

Easy override, rate, and improve loop

Authentic AI. Built for business outcomes.

Six principles that govern every agentic engagement we deliver — so AI moves business KPIs, not just dashboards.

01



Business problems first

Every engagement starts with a sharply defined pain point — not an AI capability looking for a use case.

02



Outcome-focused, always

Success is measured in cycle time saved, working capital released, decisions made better — not models deployed.

03



Readiness before rollout

Data, process, systems, organisation, governance — all five pillars assessed and gaps closed before agents go live.

04



Authentic intelligence, no theatre

No demoware. Every agent is grounded in your data, your rules and your guardrails — built to run, not just to show.

05



Human + agent partnership

We design augmentation, not replacement. People own judgement; agents own scale, consistency and speed.

06



Pilot, measure, scale

Real workflows, real users, fast feedback. Each agent earns the right to scale by proving outcomes in production.

A pragmatic 90-day starting roadmap

Days 1 - 30

Diagnose & shortlist

- Use-case inventory across functions
- Score by impact, feasibility, owner readiness
- Pick 2-3 lighthouse processes
- Confirm data, system and policy access

Days 31 - 60

Build & pilot

- Stand up agent platform and guardrails
- Configure and evaluate against test cases
- Run pilot with human-in-the-loop
- Capture KPIs vs baseline

Days 61 - 90

Scale & operate

- Move from in-the-loop to on-the-loop
- Roll out to full process scope
- Stand up monitoring and feedback loop
- Approve next wave of use cases

Outcome at day 90: 2-3 production agents, baselined ROI, a governance model in operation, and a costed plan for the next twelve.

WHY HOFT & AITHENTIC

We bring CFO-grade rigour to the agentic transformation journey

01

Authentic AI, augmented outcomes

We engineer agentic systems for measurable business impact — not demos.

02

Office of the CFO depth

Two decades of finance transformation experience underpins every solution.

03

Governance from day one

Risk, audit and human oversight are designed in, not bolted on later.

Let's identify your three highest-impact agentic use cases — together.

AITHENTIC

HOFT&
HUB OF FINANCE TRANSFORMATION