ELEVAND TECHNOLOGIES

AI Operations Assessment

Support Intelligence Pilot for Northstar Markets

CLIENT	**Northstar Markets (fictional)**
SECTOR	**B2B capital-markets software**
WORKFLOW	**Customer support triage and incident context**
ASSESSMENT	**Synthetic demonstration**
DATE	**August 2026**

SYNTHETIC SAMPLE

All organizations, records, metrics, and findings in this document are fictional. This sample demonstrates an assessment format and does not represent a client engagement or performance claim.

01 / EXECUTIVE SUMMARY

A controlled pilot is justified

Northstar's support team has a measurable context-retrieval bottleneck. A narrowly scoped assistant can reduce non-value-added search time while keeping customer communications and production actions under human control.

18 min	31%	14 wk
Median time spent gathering context	Tickets escalated before complete triage	Estimated pilot payback under base assumptions

Recommendation

- Proceed with a 60-day internal-assist pilot for one support queue.
- Limit retrieval to approved knowledge, incident, and customer-entitlement sources.
- Require an analyst to approve every customer-facing response and operational escalation.
- Do not permit the system to execute trades, change entitlements, or modify production systems.
- Expand only if quality, adoption, and control thresholds are met for four consecutive weeks.

Decision requested

Decision	Owner	Required by
Approve pilot discovery and security design	COO + CTO	Week 0
Authorize synthetic and sanitized evaluation dataset	Support + Security	Week 1
Approve production shadow mode	Risk owner	Week 5

02 / CURRENT STATE

Context is fragmented across six systems

Analysts manually reconstruct customer, product, and operational context before they can decide whether to answer, investigate, or escalate.

Step	Current activity	Median	Primary friction
1	Read and categorize ticket	3 min	Inconsistent category and urgency
2	Find customer and entitlement context	6 min	Multiple identifiers and systems
3	Search knowledge and past incidents	9 min	Keyword search misses relevant records
4	Draft response or escalation	8 min	Quality depends on individual experience
5	Review and send	4 min	Reviewer repeats context gathering

Observed constraints

DATA	Customer records cannot be mixed across tenants. Ticket text can contain personal or account information.
CONTROL	Support analysts may recommend actions but cannot grant entitlements or make production changes.
QUALITY	Approved documentation is incomplete and some historical incident notes are stale.
OPERATIONS	The workflow must degrade safely when the model or retrieval service is unavailable.

03 / OPPORTUNITY

Assist the decision; do not automate authority

The proposed system prepares evidence-backed recommendations inside the existing support workflow. It does not communicate with customers or change systems without an authorized person.

Capability	Pilot	Deferred	Prohibited
Ticket classification	Recommend category and urgency	Automatic routing	Closing tickets without review
Knowledge retrieval	Approved, entitled sources	Broader historical corpus	Cross-customer retrieval
Response support	Draft for analyst approval	Low-risk auto-send	Investment or trading advice
Operations	Recommend runbook and escalation	Create incident record	Production changes or trades

Target workflow

1 **INGEST** Ticket and customer scope	2 **RETRIEVE** Approved evidence only	3 **RECOMMEND** Category, urgency, draft	4 **APPROVE** Authorized analyst	5 **RECORD** Decision and evidence

Assumptions to validate

- At least 60% of tickets require information already present in approved internal sources.
- Customer identity can be mapped reliably to an entitlement scope before retrieval.
- Analysts will review recommendations within their existing ticket interface.
- Support leadership can label a representative evaluation set of at least 300 tickets.

04 / VALUE CASE

The base case depends on adoption and saved review time

This is a transparent decision model, not a performance guarantee. All inputs are synthetic and should be replaced with observed client data.

Input	Base assumption	Sensitivity
Eligible tickets per month	2,400	1,600 to 3,200
Minutes saved per eligible ticket	8	4 to 12
Analyst adoption	70%	40% to 85%
Loaded analyst cost	CAD 62/hour	CAD 50 to 80
Monthly platform and support cost	CAD 4,500	CAD 3,000 to 7,500
Pilot implementation	CAD 28,000	Fixed for model

Conservative	Base	Upside
CAD 1.5K	**CAD 9.4K**	**CAD 25.4K**
Estimated monthly net value	Estimated monthly net value	Estimated monthly net value
No payback within year one	Approx. 14-week payback	Approx. 5-week payback

Interpretation

The conservative scenario does not justify implementation on labour savings alone. The pilot should therefore include a stop condition: if accepted recommendations do not save at least six minutes on 60% of eligible tickets by week eight, pause expansion and redesign or discontinue the workflow.

05 / RISK AND CONTROL DESIGN

Controls are part of the product

The pilot should fail closed: unavailable identity, uncertain customer scope, missing evidence, or low confidence routes the ticket to the current manual process.

Risk	Rating	Required control	Evidence
Cross-customer data exposure	HIGH	Tenant scope applied before retrieval; deny on missing identity	Automated isolation tests
Unsupported recommendation	HIGH	Citations required; approved sources; abstain below threshold	Groundedness evaluation
Unauthorized action	HIGH	No production tools; analyst approval; least privilege	IAM and tool inventory
Sensitive data in logs	MED	Redaction; encryption; restricted retention	Log sample and retention test
Stale knowledge	MED	Owner and expiry metadata; freshness ranking	Source-age dashboard
Model/service outage	LOW	Manual workflow remains available; timeout and circuit breaker	Resilience exercise

Data handling boundary

ALLOWED IN PILOT	EXCLUDED FROM PILOT
Sanitized support tickets Approved product knowledge Customer entitlement identifiers Incident summaries Non-secret operational metadata	Credentials and secrets Payment card data Trading instructions Material non-public information Unapproved customer exports

06 / REFERENCE ARCHITECTURE

AWS-native, observable, and replaceable

The design separates orchestration, retrieval, model access, policy, and audit so components can be evaluated and replaced independently.

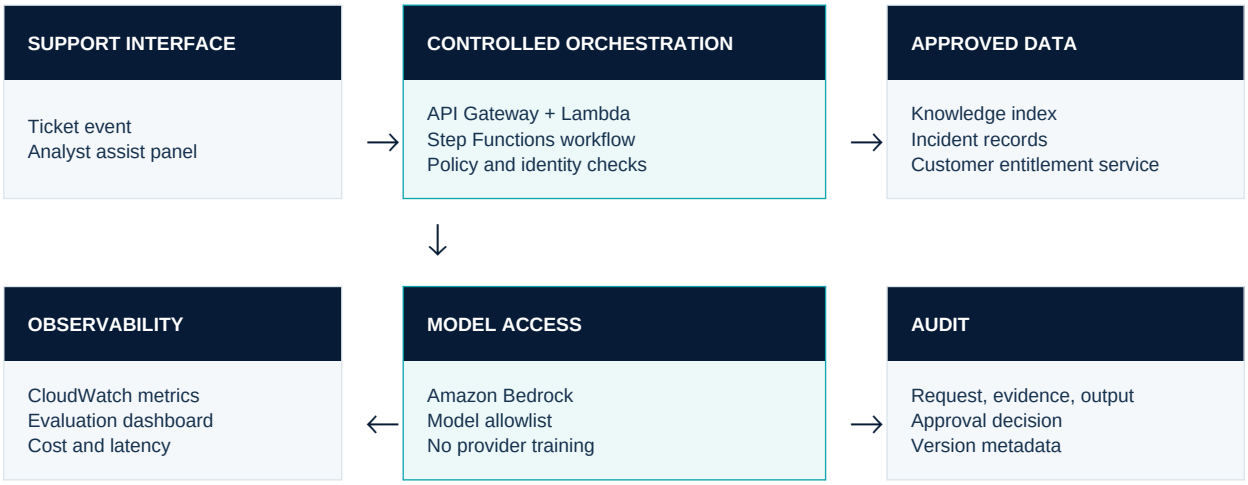

Architecture decisions

Decision	Rationale
Identity before retrieval	Reduces cross-customer leakage risk and simplifies evidence
Retrieval citations required	Makes analyst verification and audit practical
No agent write tools in pilot	Limits blast radius while quality and adoption are established
Model accessed through abstraction	Enables model comparison and reduces supplier lock-in
Evaluation in delivery pipeline	Prevents unmeasured model, prompt, or knowledge changes

07 / EVALUATION PLAN

Quality gates before production use

The pilot is accepted only when it improves the workflow without weakening privacy, customer isolation, or decision authority.

Dimension	Metric	Pilot gate	Measurement
Classification	Category agreement	>= 90%	300 labelled historical tickets
Urgency	Critical-ticket recall	>= 98%	Adversarial and historical set
Retrieval	Useful evidence in top 5	>= 92%	Reviewer relevance labels
Grounding	Unsupported material claims	< 1%	Weekly sampled review
Isolation	Cross-tenant retrieval	0	Automated negative tests
Adoption	Eligible tickets using assist	>= 60%	Product telemetry
Efficiency	Median handling-time reduction	>= 20%	Matched queue comparison

Release gates

GATE 1	**Offline evaluation**	All isolation and critical-recall thresholds met
GATE 2	**Shadow mode**	Recommendations hidden from analysts; compare to actual outcomes
GATE 3	**Limited assist**	Ten trained analysts; one queue; daily quality review
GATE 4	**Pilot acceptance**	Four consecutive weeks above quality, control, and adoption gates

08 / 90-DAY ROADMAP

A staged path with explicit stop conditions

The roadmap front-loads data controls and evaluation. Production access remains out of scope throughout the pilot.

Phase	Weeks	Delivery	Exit evidence
Foundation	1-2	Workflow baseline; source owners; threat model; evaluation set	Approved design and labelled data
Prototype	3-4	Retrieval; classification; response draft; audit schema	Offline thresholds met
Shadow	5-6	Live ticket feed; recommendations hidden; daily evaluation	No isolation failures
Limited assist	7-10	Ten analysts; human approval; feedback capture	Quality and adoption trend
Decision	11-12	Value review; control evidence; operating model	Scale, redesign, or stop

Accountability

Role	Accountability
Executive sponsor	Business outcome, funding, and final scale decision
Support owner	Workflow, analyst adoption, and operating metrics
Security and risk	Data boundary, control acceptance, and incident path
Platform owner	Cloud architecture, reliability, access, and cost
Elevand	Design, implementation, evaluation, documentation, and knowledge transfer

Stop conditions

- Any cross-customer retrieval or unapproved source access.
- Critical-ticket recall below 98% after remediation cycle.
- Unsupported material claims above 1% in reviewed recommendations.
- Adoption below 40% after workflow and training adjustments.
- No credible path to positive value under observed production assumptions.

09 / PROPOSED NEXT STEP

Confirm the workflow with evidence

A real assessment replaces every synthetic assumption in this sample with observed data, named owners, documented constraints, and an agreed decision model.

Assessment deliverables

- Current-state workflow and baseline
- Ranked opportunity and control map
- Target architecture and data boundary
- Evaluation plan and acceptance gates
- Value model with sensitivity analysis
- 90-day implementation roadmap
- Executive proceed, redesign, or stop recommendation

STARTING QUESTION

Which support, platform, or knowledge workflow consumes the most expert time without creating corresponding customer value?

ELEVAND TECHNOLOGIES

AI systems for financial operations.

Toronto, Canada | hello@elevand.ca | elevand.ca