DIGITAL CAVALIER TECHNOLOGY SERVICES INC. (DCTS)

# Decision Intelligence & Trusted AI

A practical architecture direction for explainable, auditable and human-supervised AI-assisted decisions.

| CANADIAN AI TECHNOLOGY | REAL-TIME DECISIONING | PRIVACY-BY-DESIGN |
|---|---|---|

**As AI moves from generating information to influencing consequential decisions, organizations increasingly need to understand not only what an AI system recommends, but why, on what basis, under which rules and with what human accountability. DCTS is exploring this problem through a Governance & Decision Intelligence layer built on its existing real-time decisioning technology.**

## The Problem: From AI Output to Accountable Decisions

Many organizations can deploy models or generative AI tools. The harder operational question is how to govern the decision that follows. High-assurance environments require mechanisms that make AI-assisted actions reviewable and compatible with human responsibility.

- What information influenced the recommendation?
- Which rules, thresholds or models were applied?
- Can an authorized person review or override the result?
- Can the organization reconstruct the decision later?
- Can policy and governance controls be applied consistently across systems?

## A Decision Intelligence Layer

DCTS's architectural direction is not to replace every model or enterprise system. Instead, a decision-intelligence layer can sit between data/model outputs and the business or operational action, applying context, governance logic and traceability.

### 1. Signal & Context Ingestion

Receive relevant signals from applications, models, data sources or operational systems through APIs and integration interfaces.

### 2. Decisioning

Apply rules, contextual logic, scoring and applicable model outputs to determine a recommended response or priority.

### 3. Governance Controls

Apply thresholds, policy constraints, escalation requirements and human-approval rules appropriate to the use case.

### 4. Explanation & Record

Capture sufficient context to explain the basis of an AI-assisted decision and support later review or audit.

### 5. Human Oversight

Enable authorized users to review, approve, reject or override decisions where the risk profile requires human accountability.

## Why DCTS Has a Relevant Starting Point

The proposed direction builds on technology already developed for real-time personalization and decisioning rather than beginning with a blank-sheet concept.

- Real-time signal processing and decisioning heritage.
- Approximately 94,000 lines of proprietary source code.
- Approximately 58,000 engineering hours.
- API and embedded integration capability.
- Privacy-conscious design heritage based on minimizing dependence on personally identifiable information.
- Historical sub-40-millisecond response performance at the load-balancer layer.

## Trusted AI Design Principles

### Transparency

Organizations should be able to understand the relevant factors and logic surrounding an AI-assisted decision at an appropriate level for the user and risk.

### Auditability

Decision context and governance events should be captured in a manner that supports review, investigation and assurance.

### Human Oversight

Automation should not eliminate accountable human authority where policy, risk or mission requirements demand review.

### Governance by Design

Controls should be part of the operating architecture rather than a reporting exercise added after deployment.

### Integration over Replacement

A governance and decisioning layer should be capable of working with existing models, applications and enterprise technology wherever practical.

## Potential High-Assurance Applications

The architecture is intentionally application-agnostic. Potential areas for validation include financial services, regulated enterprise environments, public sector, critical infrastructure, defence and dual-use systems.

- AI-assisted operational decision support.
- Prioritization and information triage.
- Policy-aware recommendations and escalation.
- Explainable automated workflows.
- Human-in-the-loop AI operations.
- Decision records for audit, assurance and after-action review.

## Path to Validation

DCTS's preferred approach is to validate a specific operational problem with a customer or strategic partner, identify the minimum decision-governance requirements, and then determine which existing technology components can be reused or adapted. This reduces speculative development and aligns investment with a validated need.

Important positioning note: References to defence, dual-use and high-assurance applications describe potential applications of DCTS's existing technology foundation. They should not be interpreted as claims of current defence deployment, certification or security clearance.