THE AI CONTROL PLANE

FOR FINANCIAL SERVICES.

GOVERN AI. CONTROL AGENTS. PROVE TRUST.

From AI development to autonomous execution, CognitiveView helps financial institutions build, test, govern and control AI across the enterprise. Unify AI governance, risk management, AI assurance, business-defined policies and agent runtime controls in one integrated platform.

BOOK A 30-MINUTE MEETING	MEET US AT THE AUSTRade PAVILION
18–20 November 2026 Singapore EXPO	Hall 5 | Booth 5A19

LATEST ANNOUNCEMENTS

ANNOUNCEMENT 01 | TECHNOLOGY INNOVATION

CognitiveView now supports NVIDIA OpenShell

BRINGING ENTERPRISE GOVERNANCE CLOSER TO AUTONOMOUS AI AGENT RUNTIMES.

AI agents are moving beyond generating responses to executing actions, interacting with tools and operating autonomously across enterprise environments. CognitiveView extends its AI Control Plane capabilities to supported NVIDIA OpenShell runtime environments, connecting enterprise governance requirements with agent runtime controls. This enables organisations to establish business-defined policies, govern agent permissions, define operational boundaries and extend governance into supported runtime environments. By combining governance, runtime controls and continuous assurance, CognitiveView helps enterprises move towards more controlled, accountable and trustworthy autonomous AI operations.

What this means for enterprises

- Govern agent identities, mandates and permissions.
- Define business policies and operational boundaries.
- Extend governance into supported runtime environments.
- Establish safeguards for autonomous agent activities.
- Monitor agent decisions, exceptions and control effectiveness.
- Maintain traceable governance and audit-ready evidence.

YOUR AI AGENTS. YOUR BUSINESS RULES. GOVERNED EXECUTION.

SINGAPORE FINANCIAL SERVICES

MAS SAFR | FEATURED ANNOUNCEMENT

ANNOUNCEMENT 02 | SINGAPORE FINANCIAL SERVICES

Operationalising SAFR principles with the CognitiveView AI Control Plane

GOVERN EVERY AGENT ACTION. AUTHORISE EVERY DECISION. CAPTURE EVERY OUTCOME.

As financial institutions embrace autonomous AI agents capable of initiating payments, executing trades, supporting credit decisions and performing financial operations, governance must extend beyond model assessment and policy documentation to the point where agents actually act. CognitiveView helps financial institutions operationalise the principles of MAS SAFR (Safeguards for Agentic Finance at Runtime) through a unified AI governance and assurance platform. Our capabilities support the establishment of agent identity and authority, business-defined controls, delegated mandates, runtime governance decisions, human escalation and traceable audit evidence. By connecting enterprise AI governance with supported agent runtime environments, CognitiveView helps financial institutions establish a structured approach to governing autonomous agent activities across their operations.

KEY GOVERNANCE CAPABILITIES

- Discover and register AI agents and their owners.
- Define agent identity, mandates and delegated authority.
- Establish financial limits and operational boundaries.
- Govern proposed agent actions against applicable controls.
- Support AUTO-EXECUTE, OBSERVE, ESCALATE and DENY workflows.
- Establish human oversight and approval requirements.
- Monitor exceptions and governance effectiveness.
- Maintain structured governance and audit evidence.

FROM AGENT AUTONOMY TO ACCOUNTABLE FINANCIAL OPERATIONS.

SINGAPORE AI GOVERNANCE ECOSYSTEM

AI VERIFY | RESPONSIBLE AI ASSURANCE

ANNOUNCEMENT 03 | SINGAPORE AI ECOSYSTEM

Advancing responsible AI assurance in Singapore

TEST. EVALUATE. VALIDATE. DEMONSTRATE TRUST.

As Singapore advances its responsible AI testing and governance ecosystem through AI Verify, organisations need practical ways to evaluate AI systems, validate safeguards and demonstrate responsible deployment. CognitiveView supports organisations seeking to align their AI assurance activities with the Singapore AI Verify ecosystem. Through AI testing and evaluation, risk assessment, governance controls and structured evidence, CognitiveView helps organisations assess AI quality, safety, fairness, reliability and transparency. Our platform connects evaluation outcomes with governance workflows, AI readiness assessments and continuous assurance, helping organisations move from responsible AI principles to measurable operational practices.

KEY AI ASSURANCE CAPABILITIES

- Assess AI risks and deployment readiness.
- Evaluate quality, safety, fairness and reliability.
- Validate safeguards and identify governance gaps.
- Document evaluation results and assurance evidence.

- Generate structured AI readiness reports.
- Monitor AI systems and reassess material changes.

TURN RESPONSIBLE AI PRINCIPLES INTO DEMONSTRABLE ASSURANCE OUTCOMES.

THE COGNITIVEVIEW PLATFORM

ONE PLATFORM. COMPLETE AI GOVERNANCE.

Six connected capabilities to govern and assure AI throughout the enterprise lifecycle.

01 | DISCOVER & INVENTORY Identify AI assets, models, agents and use cases. Establish visibility across the enterprise AI landscape.	**02 | GOVERN & CONTROL** Define business policies, risk requirements, permissions and operational guardrails.
03 | TEST & EVALUATE Evaluate AI quality, safety, fairness, security and reliability.	**04 | AGENT & RUNTIME CONTROL** Govern agent identity, mandates, permissions and safeguards across supported runtime environments.
05 | MONITOR & MANAGE RISK Monitor AI behaviour, incidents, exceptions, risk exposure and control effectiveness.	**06 | AUDIT & ASSURE** Generate traceable evidence, governance documentation and audit-ready reports.

GOVERN THE COMPLETE AI LIFECYCLE

BUILD — Connect developer workflows, AI applications, models, agents and infrastructure.

VALIDATE — Assess AI risks, evaluate performance, test safeguards and validate readiness.

DEPLOY — Establish governance approvals, policy requirements and deployment readiness.

OPERATE — Govern agent activities, monitor decisions, manage exceptions and enforce supported controls.

ASSURE — Maintain structured evidence, audit records and continuous governance feedback.

SAFR IN ACTION

FROM AGENT DECISION TO SETTLEMENT — GOVERNED, ASSURED AND AUDITABLE.

How can financial institutions govern an autonomous agent before it executes a consequential action?

1. PROPOSED AGENT ACTION Action details, context and reasoning.	**2. AGENT IDENTITY** Registered agent, ownership and mandate.
3. CONTROLS REPOSITORY Policies, limits, authority and risk thresholds.	**4. DISPOSITION ENGINE** Evaluate against controls and determine outcome.

GOVERNANCE OUTCOMES

AUTO-EXECUTE Permit within approved limits.	**OBSERVE** Execute with monitoring.
ESCALATE Route to human review.	**DENY** Block out-of-policy actions.

AUDIT LOG & GOVERNANCE EVIDENCE — Decision, rules applied, outcome and full traceability.

FINANCIAL SERVICES

BUILT FOR FINANCIAL SERVICES

Govern autonomous AI across critical financial operations.

PAYMENTS & TREASURY Transaction limits, delegated authority, agent permissions and operational controls.	**LENDING & CREDIT** Decision boundaries, risk limits, approval workflows and human oversight.
WEALTH & ADVISORY Agent mandates, permissions and decision accountability.	**COMPLIANCE & OPERATIONS** Governance checks, workflows, exceptions and audit records.

CONNECTED ENTERPRISE ECOSYSTEM

Developer & AI Infrastructure
GitHub | GitLab | Jenkins | CI/CD Pipelines | APIs & SDKs | Model Registries

AI Testing & Evaluation
DeepEval | AI Verify | Promptfoo | Evidently AI | Arize Phoenix | MLflow

Enterprise GRC / ERM
Archer | ServiceNow | IBM OpenPages | MetricStream | OneTrust | Other GRC Systems

Runtime Environments
NVIDIA OpenShell | Microsoft Agent Governance Toolkit

Standards & Regulatory Alignment
MAS SAFR | NIST AI RMF | ISO/IEC 42001 | EU AI Act | OECD AI Principles | Sector Guidance

FIND US AT THE SINGAPORE FINTECH FESTIVAL



AUSTRade PAVILION

Hall 5 | Booth 5A19

18–20 November 2026

Singapore EXPO

Use the sample pavilion image to help identify the Team Australia stand.

LET'S BUILD THE NEXT GENERATION OF TRUSTED FINANCE.

Discuss your AI governance, SAFR readiness or autonomous agent use case with our team.

BOOK A 30-MINUTE MEETING

https://calendly.com/dilip-cognitiveview/30-minute-introductory

COGNITIVEVIEW | The AI Control Plane for Trusted AI and Autonomous Agents.

Website: https://www.cognitiveview.com | Email: sales@cognitiveview.com

CONTROL AI. PROVE IT.

Positioning note: SAFR is described here as an industry reference approach that CognitiveView helps operationalise. AI Verify alignment is not presented as certification or endorsement. NVIDIA OpenShell support refers to supported runtime environments.