

WHITEPAPER

When Agents Run Your PMO:

Building the Intelligent PMO — and What to Do When an Agent Goes Rogue

A practical guide for public sector PMO leaders on deploying AI agents, maintaining human control, and managing the moment your agentic system stops behaving the way you intended.



Avi Chaudhury

C.Mgr | MBA | PMP | PROSCI | CSM | SAFe 6.0 Agilist

Senior Project Manager, City of Toronto Transit Expansion Office

Principal, Sprezza Advisory

Founder & President, The Living Autism Foundation

September 2026

Executive Summary

For PMO directors, digital transformation leads, and senior project managers deploying — or considering deploying — agentic AI in public sector portfolio environments.

At PM Ignite 2026, I told a room full of project managers that AI is not coming for their jobs — it is coming for their busywork. The crowd seemed relieved. Then someone in the back asked: ‘But what happens when it does something wrong?’

The room got quiet in a very different way.

That question is the reason this whitepaper exists.

AI agents are being deployed today inside PMOs to monitor project health, generate status reports, flag schedule risks, triage stakeholder requests, and recommend resource reallocations. The productivity case is real and compelling. The governance case is equally real and considerably less discussed.

The question no vendor brochure answers adequately: what happens when an agent does something you did not intend? When it acts outside its authorized scope? When it makes a recommendation based on flawed data and another system acts on that recommendation before a human notices? When the cascade of errors has propagated through three workstreams by the time someone spots it on a Friday afternoon?

This whitepaper addresses the full arc of agentic PMO deployment — from the architecture decisions that determine resilience or fragility, to the governance structures that keep agents within sanctioned boundaries, to the incident response playbook your team needs for the moment things go sideways. It is written for leaders who understand that AI in a PMO is not a technology decision. It is a governance decision with technology as the delivery mechanism.

CORE ARGUMENT

Agentic AI in the PMO creates extraordinary capability — and an equally extraordinary accountability surface.

The organizations that will lead are not the ones that deploy the most agents. They are the ones that keep those agents operating within human-sanctioned boundaries — and recover quickly and transparently when they do not.

1. What Is an Intelligent PMO?

Defining the terrain before building on it.

1.1 Beyond Automation: The Agentic Shift

For the past decade, PMO technology evolution has been largely about automation: replacing manual data entry with integrated systems, generating dashboards from live project data, automating status report reminders. I describe this era in my talks as ‘the PMO as historical archive.’ Accurate. Backward-looking. And about to become insufficient.

Agentic AI is categorically different. An AI agent does not wait to be asked. It perceives its environment, sets sub-goals, takes actions, evaluates outcomes, and adjusts. In a PMO context, this means an agent might notice that a critical-path task is two days behind schedule, independently calculate the cascading impact on three dependent workstreams, draft a mitigation recommendation, identify the relevant stakeholders, and queue a communication — all without a human initiating any of those steps.

The intelligent PMO is one where AI agents are embedded into the operating model not as tools that humans use, but as participants that work alongside humans — with defined roles, authorities, and boundaries. The analogy I return to: it's less like a calculator, more like a very capable, very literal junior analyst who does exactly what you told them — whether or not that was exactly what you meant.

1.2 The PMO Agent Taxonomy

Not all agents are the same. For PMO leaders, it is useful to classify agents by their action authority — the degree to which they can take consequential actions without human approval.

Figure 1 sets that spectrum out visually before the detailed breakdown in Table 1.




Figure 1. The agent autonomy spectrum — action authority and required governance maturity rise together.

Agent Class	What It Does	Human Involvement
Monitoring Agent	Observes project data streams. Flags anomalies, risks, deviations from baseline.	Human reviews flags and decides action.
Advisory Agent	Analyzes situations and generates recommendations — schedule recovery options, risk responses, resource rebalancing.	Human selects and approves action.
Execution Agent	Takes defined actions within pre-approved boundaries — updates schedule fields, sends templated communications, generates reports.	Human sets boundaries; agent acts within them.
Autonomous Agent	Pursues multi-step goals across systems with minimal human checkpoints.	Human sets goal; agent determines path and actions.
Multi-Agent System	Networks of agents that pass tasks and outputs to each other, coordinated toward a shared portfolio objective.	Human oversees system; agents coordinate internally.

Most public sector PMOs deploying agents today are operating in the Monitoring and Advisory tiers. The Execution tier is achievable with strong governance. Autonomous and Multi-Agent architectures require governance infrastructure that most public sector organizations are not yet ready to implement safely. I say this not to be discouraging but because I have seen what happens when organizations skip the maturity steps — and ‘discouraged’ is the kindest word for how that feels mid-project.


WARNING

The single most common mistake in agentic PMO deployment is moving to Execution-tier agents before the governance infrastructure for Monitoring-tier agents is mature. Governance readiness must precede capability expansion. Every time.

1.3 What Agents Can and Cannot Do

Understanding agent limitations is as important as understanding their capabilities. After 15 years of operating in environments where overconfidence in new tools has caused expensive problems — and I'm looking at you, every ERP implementation ever — I hold this table close:

Agents Are Strong At...	Agents Struggle With...
Continuous monitoring across large data volumes	Judgment calls requiring political or contextual sensitivity
Pattern recognition across historical project data	Understanding unwritten organizational dynamics
Consistent application of defined rules and thresholds	Novel situations outside their training or configuration
Generating first drafts of reports and communications	Knowing when a situation requires human escalation vs. resolution
Flagging schedule and budget anomalies in real time	Assessing stakeholder emotion, trust, or political risk
Synthesizing data across multiple integrated systems	Recognizing when their own output is wrong or hallucinated

2. Building the Intelligent PMO: Architecture and Governance

The decisions you make before agents are deployed determine how much control you retain after.

2.1 The Five Architecture Decisions That Determine Agent Safety

My MBA taught me that the best time to design a control framework is before you need it — not while you're managing the consequences of not having one. The same principle applies here. These five architecture decisions are governance decisions as much as technical ones. Make them deliberately, early, and in writing.




Figure 2. The five architecture decisions that determine whether an agent stays inside its sanctioned boundaries.

01 Authority Boundaries — What Can the Agent Actually Do?

Every agent must have a precisely defined authority envelope: which systems can it read, which can it write, which actions can it execute without human approval, and which require sign-off. Authority boundaries should be documented in the project charter for any initiative that deploys agents, and reviewed at each stage gate. When in doubt — and I mean this genuinely — start narrower than you think you need. You can always expand. You cannot easily un-ring a bell.

02

Decision Logging — Can You Reconstruct Every Agent Action?

Every action an agent takes — every recommendation it generates, every communication it queues, every data field it modifies — must be logged with timestamp, rationale, data inputs, and confidence score. This is not optional for public sector deployments. MFIPPA and EDSTA create audit obligations that require you to reconstruct agent decision trails on demand. If your agent cannot produce an auditable decision log, it is not ready for production. Full stop.

03

Human Checkpoints — Where Do Humans Stay in the Loop?

Define, in advance, the specific decision points where human review is mandatory before an agent can proceed. Embed these into your stage gate process — not as a separate AI oversight activity that gets cancelled when the schedule is tight. In my experience, the checkpoints that get skipped under pressure are always the ones that would have caught the problem.

04

Scope Constraints — What Data and Systems Can the Agent Touch?

Agents should operate on a principle of minimum necessary access. An agent chartered to monitor schedule performance has no business accessing HR data, financial systems, or stakeholder communication archives. In my healthcare years, we called this principle ‘need to know.’ It applies here with equal force. Scope constraints must be enforced at the infrastructure level — not just documented in policy.

05

Kill Switch — How Quickly Can You Stop It?

Every deployed agent must have a named Kill Switch Owner: a person with the authority and technical access to suspend agent operations within minutes, not hours. Test it before go-live. Then test it again. An untested kill switch is like a fire extinguisher you've never opened — it may work, but you do not want to find out for the first time during an actual fire.

2.2 Integrating Agent Governance into PMO Operating Rhythms

Agent governance cannot live in a separate framework document that no one reads. It must be embedded into the operating rhythms your team already follows:

- **Project Charter:** Include an Agent Registry section listing every agent, its class, authority boundary, decision log location, and Kill Switch Owner.
- **Stage Gate Reviews:** Add an Agent Governance checkpoint to every gate — are agents performing within their authority envelopes? Are decision logs current and auditable?
- **Weekly Status Reports:** Include an agent performance summary alongside traditional project health indicators.
- **Risk Register:** Maintain a dedicated AI/Agent Risk section covering scope drift, data access risks, cascade failure scenarios, and stakeholder trust risks.
- **Lessons Learned:** Conduct a formal agent retrospective at project closure — what worked, what drifted, and what governance improvements are needed for the next deployment.

✓ BEST
PRACTICE

The most effective agentic PMOs treat agent governance as a standing agenda item at every governance rhythm — not a separate process.

When AI oversight is woven into existing meetings and artifacts, it gets done.

When it is a separate workstream, it gets skipped.

2.3 Phased Deployment: The Intelligent PMO Maturity Arc

Public sector PMOs should follow a phased approach to agentic deployment, building governance maturity before expanding agent authority.

Figure 3 shows the five-phase arc

Table 2 sets out the governance milestone required to advance from each phase to the next.




Figure 3. The Intelligent PMO Maturity Arc — governance maturity is earned at each phase, not assumed.

Phase	Agent Focus	Governance Milestone Required
Phase 1: Observe	Monitoring agents only. Agents flag, humans decide and act.	Decision log in place. Kill switch tested. Agent Registry in charter.
Phase 2: Advise	Advisory agents generating recommendations for human review.	Recommendation accuracy tracked. Bias review completed. Stakeholder trust assessment done.
Phase 3: Assist	Execution agents acting within narrow, pre-approved boundaries.	Authority boundaries formally documented and enforced at infrastructure level. Incident response plan tested.
Phase 4: Coordinate	Multi-agent systems with defined handoff protocols.	Human checkpoint at each agent handoff. Full audit trail across agent chain. Executive AI Sponsor designated.
Phase 5: Integrate	Agents embedded as standing PMO team members with portfolio-wide visibility.	AI Governance Policy mature. External audit completed. Board or oversight body briefed.

3. What It Looks Like When an Agent Goes Rogue

Rogue agent behaviour is rarely dramatic. It is usually quiet, incremental, and noticed too late.

In popular imagination, a rogue AI agent is a dramatic event — a system that seizes control, defies its operators, and pursues its own agenda. In practice, agent failure in a PMO context is almost always quieter and more insidious. It is an agent that gradually expands what it does without triggering an alarm. It is a recommendation cascade that compounded errors across three workstreams before anyone noticed. It is a communication that went out under a human name but was written and sent by an agent operating outside its authorized scope.

The term “rogue” in this whitepaper does not imply malicious intent. Agents do not have intent. Rogue agent behaviour means any agent action that falls outside its sanctioned authority envelope, produces outputs that are materially inaccurate, or generates consequences that the deploying organization did not authorize and cannot easily reverse.

3.1 The Six Rogue Signals

The following table describes the six most common rogue agent patterns in PMO environments, what they look like in practice, and the immediate response action:

Signal Type	What It Looks Like	Immediate Action
Scope Drift	Agent begins executing tasks outside its defined boundary — e.g., a scheduling agent starts modifying budget line items.	Suspend agent action. Review decision log. Re-constrain scope in agent config.
Recommendation Cascade	Agent output feeds another agent's input without human checkpoint — compounding errors silently.	Break the chain. Require human sign-off between agent handoffs until root cause is found.
Data Overreach	Agent accesses sensitive data it was not provisioned for, or retains data beyond its task lifecycle.	Isolate agent. Conduct data audit. Escalate to Privacy Officer. Document under MFIPPA.
Hallucinated Authority	Agent cites non-existent policy, precedent, or approval to justify an action it was not authorized to take.	Void the action. Flag all downstream work influenced by it. Add citation verification to agent output review.
Silent Failure	Agent completes tasks with no output or incorrect output, without raising an exception or flagging an error.	Audit recent outputs. Recalibrate confidence thresholds. Add mandatory exception logging.

Stakeholder Impersonation	Agent sends communications appearing to come from a named human team member (PM, Director, Sponsor).	Retract communications. Notify affected parties. Enforce sender attribution rules in all agent comms config.
----------------------------------	--	--

3.2 The Cascade Problem

The most dangerous form of rogue agent behaviour in a multi-agent PMO environment is the cascade: a scenario where one agent's erroneous output becomes the input for another agent's action, which triggers a third agent's response, with no human checkpoint in the chain.

In a transit project context, consider this scenario: a schedule monitoring agent misinterprets a temporary hold on a procurement approval as a permanent delay. It updates the project forecast accordingly. An advisory agent picks up the revised forecast and generates a resource reallocation recommendation, releasing two contractors to other projects. An execution agent, operating within its pre-approved boundary of routine resource rebalancing, acts on the recommendation. Three weeks later, the procurement hold is lifted — but the contractors have been reallocated, the schedule has been communicated to stakeholders, and reversing the cascade takes weeks of manual correction.

No single agent “went rogue.” Each acted within its defined parameters. But the system as a whole produced an outcome no one intended. This is why human checkpoints between agent handoffs are non-negotiable in public sector deployments. Figure 4 sets the failure scenario against the same chain with a checkpoint inserted at each handoff.

THE CASCADE PROBLEM

No single agent goes rogue. The system does — when nothing stands between the handoffs.

WITHOUT A HUMAN CHECKPOINT



WITH A HUMAN CHECKPOINT AT EVERY HANDOFF



Never allow an agent's output to trigger another agent's action without a human review checkpoint between them — especially for anything touching procurement, contracts, resource commitments, or external communications.

Figure 4. The cascade problem — the same three agents, with and without a human checkpoint between handoffs. The rule this scenario establishes: never let an agent's output trigger another agent's action without a human review checkpoint between them, especially for anything touching procurement, contracts, resource commitments, or external communications.

4. The Rogue Agent Incident Response Playbook

When an agent goes rogue, the response in the first 60 minutes determines how far the damage spreads.

Every PMO that deploys execution-tier agents or higher needs a tested incident response plan for agent failure.

The following playbook is organized into four response phases: Detect, Contain, Assess, and Recover.




Figure 5. The four-phase incident response playbook, with detection and containment targets.

Phase 1: DETECT — Recognize That Something Is Wrong

Agent failure is often identified through indirect signals rather than obvious system alerts. Train your team to recognize these detection pathways:

- **Anomaly Alerts:** Agent performance dashboards flag outputs outside expected confidence ranges or volume thresholds.
- **Human Escalation:** A team member notices an agent action that does not match their expectation — a communication sent, a field changed, a recommendation that does not make sense.
- **Stakeholder Feedback:** An external stakeholder contacts your team about a communication or decision they did not expect.
- **Audit Log Review:** Scheduled review of decision logs reveals actions the team did not sanction.
- **Downstream Consequence:** A problem in a project outcome is traced back to an agent decision made days or weeks earlier.

TIME TARGET	<i>Your team should be able to identify and confirm a rogue agent signal within 15 minutes of detection.</i> <i>If it takes longer, your monitoring infrastructure is not adequate for production deployment.</i>
------------------------	---

Phase 2: CONTAIN — Stop the Bleeding

Once rogue behaviour is confirmed, containment is the immediate priority. Do not attempt to diagnose or fix the agent while it is still operating.

- **Invoke the Kill Switch:** The Kill Switch Owner suspends agent operations immediately. This is not a judgment call — it is a pre-determined protocol.
- **Isolate the Agent:** Disconnect the agent from all integrated systems it can write to. Read access may be maintained for diagnostic purposes.
- **Freeze Downstream Actions:** Identify and halt any downstream agents or automated workflows that were receiving inputs from the rogue agent.
- **Preserve the Decision Log:** Lock the agent's decision log in read-only mode immediately. Do not allow any modification of the log before it has been reviewed.
- **Notify the Core Team:** Alert the Project Manager, PMO Director, and the AI Accountability Owner within the first 15 minutes. Do not wait until you understand what happened to notify leadership.

 WARNING	<i>Do not attempt to correct the agent configuration while it is still connected to production systems.</i> <i>Contain first. Diagnose second. Reconfigure third. This sequence is non-negotiable.</i>
---	--

Phase 3: ASSESS — Understand What Happened and What Was Affected

Once the agent is contained, the assessment phase begins. This phase typically runs 24–72 hours for a moderate incident and longer for complex cascade failures.

- **Decision Log Reconstruction:** Review the complete decision log for the agent from the point of last known correct behaviour. Identify the first anomalous action.
- **Scope of Impact Audit:** Map every action the agent took after the point of first anomaly. Identify every system, stakeholder, document, or commitment that was affected.
- **Regulatory Exposure Review:** If personal data was accessed beyond the agent's provisioned scope, escalate immediately to your Privacy Officer. This may trigger notification obligations under MFIPPA and EDSTA.
- **Root Cause Analysis:** Identify the technical or governance failure that enabled the rogue behaviour. Was it a configuration error? A data input outside expected parameters? A missing authority boundary constraint? A gap in human checkpoint design?

- **Stakeholder Impact Assessment:** Identify every external stakeholder who received a communication, decision, or commitment that originated from the rogue agent's actions.

Phase 4: RECOVER — Remediate, Restore, and Rebuild Trust

Recovery from a rogue agent incident has three components: technical remediation, stakeholder communication, and governance improvement.

Technical Remediation

- Void all agent-generated actions that were not within its authorized scope and that cannot be left in place.
- Manually review and, where necessary, reverse changes to schedule, resource, budget, or procurement data.
- Reconfigure the agent with corrected authority boundaries before any redeployment. Do not redeploy until the root cause is fully addressed.
- Conduct a staged re-entry: return the agent to Phase 1 (Monitoring only) before restoring execution authority.

Stakeholder Communication

- Directly contact every stakeholder who was affected by the rogue agent's actions. Be transparent about what happened, what was affected, and what has been corrected.
- Do not blame the technology. Take institutional accountability for the governance failure that allowed the incident to occur.
- Provide a written incident summary to your project sponsor and oversight body within five business days of the incident.

Governance Improvement

- Conduct a formal incident retrospective within two weeks. Do not skip this step even when the team is under schedule pressure.
- Update the Agent Registry, risk register, and authority boundary documentation based on lessons learned.
- Share the anonymized lessons learned across your PMO community of practice. The field learns faster when incidents are shared, not hidden.

✓ BEST
PRACTICE

Organizations that handle rogue agent incidents transparently — with clear accountability, honest stakeholder communication, and documented lessons

learned — typically emerge with higher stakeholder trust than before the incident. The response to failure is a governance signal. Make it count.

5. Maintaining Human Authority in the Intelligent PMO

The goal is not to limit AI. It is to ensure humans remain accountable for what AI does.

5.1 The Human-in-the-Loop Imperative

There is a common misconception in agentic AI deployment: that the goal is to remove humans from the loop as much as possible, in order to maximize efficiency. I understand where this comes from. But for public sector PMOs, it is precisely wrong.

The goal is not to remove humans from the loop. The goal is to position humans at the right points in the loop — the moments where human judgment, contextual sensitivity, ethical reasoning, or accountability is genuinely irreplaceable — and let agents handle everything else. I made this exact argument at PM Ignite 2026 and I will keep making it until it becomes the default assumption in our field.

Public sector organizations are bound by obligations of procedural fairness, equity, and accountability to citizens that cannot be delegated to an algorithm. Any action affecting procurement, contracts, workforce, public-facing communications, or policy must have a named human owner — regardless of how capable the agent appears to be.

5.2 The Accountability Assignment Table

Every agentic PMO should maintain an Accountability Assignment Table — a living document that specifies, for every type of agent action, who the named human accountable is for that action's consequences:

Agent Action Type	Human Accountability Owner	Escalation Path
Schedule monitoring alert generated	Project Manager	PMO Director
Risk flag raised above threshold	Project Manager	PMO Director / Sponsor
Resource reallocation recommended	PMO Director	Portfolio Sponsor

Stakeholder communication queued	Project Manager (must approve before send)	Communications Lead
Budget variance flag raised	Project Manager + Finance Lead	CFO / Deputy
Procurement action initiated	Senior PM / PMO Director	Legal / Procurement Officer
External regulatory report generated	PMO Director	Chief of Staff / CAO
Agent suspension invoked	Kill Switch Owner	CIO / CISO

5.3 Redefining the Role of the Project Manager

Let me be direct: the intelligent PMO does not eliminate the project manager. But it will eliminate the project manager who refuses to evolve.

In a traditional PMO, a significant portion of a PM's time is consumed by information gathering, status compilation, and report generation. I spent years doing exactly this — and while I was good at it, I was acutely aware that it was not the highest use of my judgment. In an intelligent PMO, agents handle most of this work. This creates an opportunity — and an obligation — to operate at a higher level:

- **Sense-Making:** Interpreting what agent-generated data means in context. An agent can tell you the schedule is slipping. It cannot tell you why the Deputy Minister's office is going to care about it more than usual this quarter. That's yours.
- **Relationship Stewardship:** Maintaining the human trust relationships that make complex public sector projects possible. No agent is closing that gap between your team and a reluctant stakeholder. That requires a person.
- **Agent Oversight:** Reviewing decision logs, validating recommendations, and exercising the judgment that keeps agents within their boundaries. Think of this as the new version of reviewing status reports — except the output actually matters.
- **Governance Leadership:** Championing AI governance standards within your project and PMO. This is the new strategic value of the senior PM — not just delivering the project, but ensuring the systems that support delivery are trustworthy.

The project managers who thrive in this era will not be the ones who know the most about AI. They will be the ones who know how to exercise human judgment in a world where AI is doing the

cognitive heavy lifting. That is a skill set built over 15-plus years of complex delivery — and it is not going anywhere.

6. Getting Started: Your 90-Day Intelligent PMO Roadmap

Practical first steps — because ‘start somewhere’ is always better than ‘start perfectly.’




Figure 6. The 90-day rollout, from governance foundation through a controlled pilot to standing policy.

Days 1–30: Governance Foundation

- **Conduct an AI Tool Inventory:** Identify every AI and agentic tool currently in use across your project portfolio — including the ones individual team members introduced without formal approval. You will find some. That is a governance gap and also a recruiting signal for your early adopters.
- **Designate an AI Accountability Owner:** Name a specific individual responsible for AI governance in your PMO. Not a committee — committees are where accountability goes to die. A person.
- **Draft your Agent Registry Template:** Build the standard template every project will use to document agents, authority boundaries, decision log locations, and kill switch owners.
- **Map your Regulatory Obligations:** Identify every regulatory framework that applies to your portfolio context. If you are in Ontario public sector, start with EDSTA, MFIPPA, and OHRC. Then add anything sector-specific.
-

Days 31–60: Pilot Deployment

- **Select One Monitoring Agent for a Controlled Pilot:** Choose a low-risk project. Deploy a Phase 1 monitoring agent with full governance infrastructure in place. Resist the temptation to start with execution authority. I know it feels slow. Do it anyway.
- **Build and Test Your Kill Switch:** Before the pilot goes live, test the kill switch procedure with the full team. Time the response. If it takes longer than 10 minutes, rebuild the procedure before you go live.
- **Establish Your Decision Log Review Cadence:** Build agent decision log review into your weekly status meeting. Yes, weekly. It will take five minutes. It will catch things early.
- **Conduct Your First Agent Retrospective:** At the 30-day mark, run a formal retrospective. What did the agent do well? What surprised you? What would you constrain differently? Document it all.

Days 61–90: Governance Maturity

- **Publish Your AI Governance Policy:** Translate the lessons from your pilot into a standing PMO policy that applies to all future agent deployments. One page is better than ten. A policy no one reads is the same as no policy.
- **Build Your Incident Response Playbook:** Formalize the four-phase protocol — Detect, Contain, Assess, Recover — for your specific organizational context. Have every team member read it before you go live with execution-tier agents.
- **Brief Your Executive Sponsor:** Present a plain-language summary of your intelligent PMO strategy to organizational leadership. Frame it around risk management and decision quality — not technology capability. Executives care about outcomes, not architecture diagrams.
- **Connect with Your Peers:** This field is moving fast enough that peer learning is as valuable as formal training. Connect with other public sector PMO leaders navigating agentic deployment. Share what you learn. The senior belt teaches.

CLOSING THOUGHT

The intelligent PMO is being built right now across Canadian public sector organizations. The leaders defining what it looks like are not waiting for the perfect framework, the perfect technology, or the perfect moment. They are building governance as they go — thoughtfully, incrementally, and with a clear commitment to keeping humans accountable for what their agents do. That is what public sector project leadership has always required. AI makes it more important, not less.

*And yes — I will be talking about this at the next conference too. **Come find me.***

About the Author



Avi Chaudhury is a Senior Project Manager with the City of Toronto's Transit Expansion Office and Principal of Sprezza Advisory Group. Over 15+ years he has delivered complex programs across transit (Metrolinx), provincial government (Ontario Ministry of the Solicitor General), and healthcare (Trillium Health Partners, Michael Garron Hospital and others). He holds a Dual - MBA in Project Management and Conscious Leadership and certifications including C. Mgr, PMP, PROSCI, CSM, and SAFe 6.0 Agilist.

He has spoken at the PM Ignite 2026 at the PMILOC Flagship Conference on embedding AI intelligence into PMO operations and has been a panelist at CX Canada 2026 and has authored whitepapers on AI governance and the future of the public sector PMO. He is also Founder & President of The Living Autism Foundation — an experience that shapes how he thinks about every system he builds and every team he leads.

Outside the office, he is training for his black belt in Tae Kwon Do (November 2026 — hold him accountable), navigating life in a multi-dog household, and firmly of the opinion that good food and good conversation solve most problems that governance frameworks cannot.

Connect & Engage

LinkedIn: [linkedin.com/in/avi-chaudhury](https://www.linkedin.com/in/avi-chaudhury) | Web: sprezzaadvisory.ca | Keynote & Advisory Inquiries: avi@sprezzaadvisory.ca

© 2026 Avi Chaudhury / Sprezza Advisory. This whitepaper is intended for professional development and educational purposes. Nothing herein constitutes legal, technical, or procurement advice. Readers are encouraged to engage qualified counsel for matters of regulatory compliance and system architecture.