FUSION COMPUTING FIELD GUIDE

RPA vs AI Agents: What Canadian SMBs Should Automate in 2026

Mike Pearlstein, CISSP, MSc Computer Science (AI)

August 30, 2026

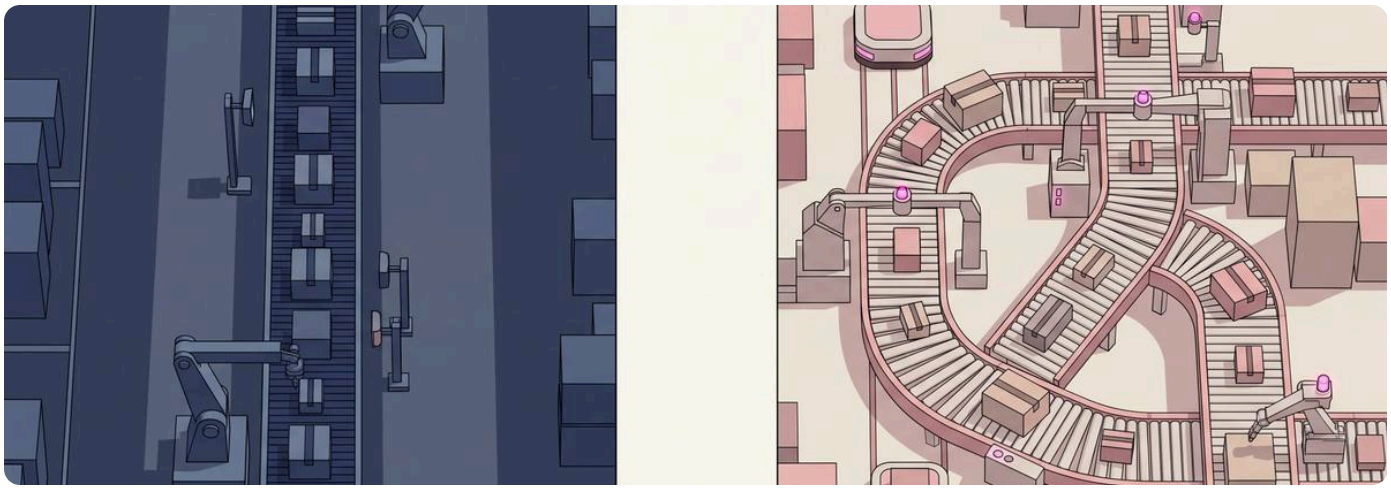

WHAT'S INSIDE

0. Frequently asked questions

TRUSTED BY **Toronto law firms** **Hamilton manufacturers** **Vancouver clinics** **GTA accounting firms** **Ontario non-profits** **British Columbia professional services**

Written by Mike Pearlstein, CISSP, CEO of Fusion Computing Limited. Helping Canadian businesses build and manage secure IT infrastructure since 2012 across Toronto, Hamilton, and Metro Vancouver.

Two different things get sold under the same label. Deterministic automation follows a script somebody wrote. Agentic automation decides its own next step. We see both sold as AI. Where the work is deterministic, we treat it as workflow automation rather than paying agent prices for a fixed path.

They fail differently, and that difference is the useful thing to know. Script failures are usually reproducible, because the same broken rule breaks the same way. Agent failures can vary between runs, which makes them harder to predict and to reproduce.

> **Short answer:** Use deterministic automation, such as robotic process automation (RPA), when the rule is fixed and the path repeats. Use an AI agent when the work needs judgement, dynamic planning, or a decision about which tool to reach for.
>
> RPA usually has a lower cost per run and is easier to audit, but is brittle when the path or interface changes. An agent can choose a different path as conditions change, and is non-deterministic, so the

same input can produce a different route. In our experience, small businesses asking for either one usually need the process written down first.

Key takeaways

- RPA is not finished. The market grew 14.5% to US$3.6 billion in 2024, though generative AI and agentic automation slowed it from 22.1% growth the year before (Gartner, 2025).
- Most “agentic” vendors are not selling agents. Gartner estimates only about 130 of the thousands of agentic AI vendors are real (Gartner, 2025).
- Agents are non-deterministic by design. One benchmark ran the same agent on the same open-source build tasks three times and scored 45.6%, 54.7%, and 58.8% (BuildBench, 2025).
- Autonomy is still limited on long-horizon work. The strongest agent finished 30.3% of 175 multi-step tasks in a simulated software company without help (TheAgentCompany, 2025).
- Gartner expects over 40% of agentic AI projects to be canceled by the end of 2027, citing cost, unclear value, and inadequate risk controls (Gartner, 2025).

What is robotic process automation, and is RPA finished?

RPA is not finished, it is growing more slowly. Gartner (2025) puts the RPA software market at US$3.6 billion in 2024, up 14.5%, down from 22.1% growth to US$3.2 billion in 2023, and attributes the slowdown to generative AI, computer use tools, and agentic automation.

Growth that decelerates is still growth, so RPA isn’t finished. Gartner doesn’t quantify RPA’s share of new automation projects, so the honest read for Canadian SMBs is category pressure rather than a measured loss

of share.

Book a Consultation

Robotic process automation is a script that drives software the way a person would. Open this, copy that, paste it there. It follows the logic you gave it, branches included, and it doesn't understand the work.

Two things follow. It's cheap to run and easy to audit, because it is deterministic: given the same input and the same screens, it produces the same output every run. It's also brittle in one specific way. Change the screen and the robot walks into a wall.

Revenue kept climbing while the growth rate fell by 7.6 points. Source: Gartner Market Share Analysis, RPA Worldwide 2024.

The "AI killed RPA" framing is wrong. RPA was often stretched onto processes with exceptions and unstructured inputs. Document AI and other workflow tools already handled some of that work before today's agent platforms, and Microsoft shipped invoice extraction inside Power Automate in 2020. Agents widen the set of cases that can adapt as they run.

What remains is what scripts were always good at. High-volume, rule-fixed steps between systems that have no API. If the process is take the number from field A and put it in field B, forever, that is still a script.

Fusion Computing builds those flows in Power Automate from CA$500 per workflow, after a free scoping call. A paid discovery engagement is CA$750 where one is needed, and RPA, AI Builder, premium connectors, or extra capacity can add licensing cost.

What does an AI agent do differently?

Assume the label on the pitch is unreliable. Gartner (2025) estimates only about 130 of the thousands of vendors selling agentic AI are real. Gartner calls this agent washing: many vendors rebranding assistants, RPA, and chatbots without substantial agentic capability. Any Canadian SMB evaluating a pitch should assume the label is unreliable and test for the one behaviour that actually distinguishes an agent.

> > *Gartner Senior Director Analyst Anushree Verma put it this way: "Most agentic AI projects right now are early stage experiments or proof of concepts that are mostly driven by hype and are often misapplied."*
>
> *Anushree Verma, Senior Director Analyst, Gartner, Gartner press release (2025)*

That's the behaviour that matters: choosing the sequence. An agent gets a goal and can choose its actions as it works, using whatever tools it is permitted. A script is told how. An agent is told what.

The capability this buys is open-ended decision making. Variable document formats on their own do not require an agent, because a deterministic flow paired with document extraction already reads a supplier PDF and a phone photo of the same invoice. Reach for an agent when the work needs

judgement, dynamic planning, or tool selection, not merely because the input varies. Agent tooling in ChatGPT put that within reach of smaller teams.

The cost is non-determinism. The same input can produce a different path on two consecutive runs. For drafting a summary that's fine. For anything writing to a financial system, it is the entire risk conversation.

Deterministic vs agentic automation, explained

The property that separates the two is repeatability. In the BuildBench study (2025), which measures agents compiling open-source software, the same agent run three times on the same tasks scored 45.6%, 54.7%, and 58.8%, a 13.2-point spread around a 53.0% average. Its success rate rose from 54.7% on one attempt to 65.5% given three, because model outputs are non-deterministic even with identical prompts. That variance is the property the comparison below turns on.

Deterministic automation (RPA) versus agentic automation, across the eight axes that decide which one a process needs.

Axis	Deterministic (RPA)	Agentic (AI agent)
What you tell it	How to do the task	What to achieve
Suits work that is	Rule-fixed on a stable path	Judgement-led, needs dynamic planning, or requires tool selection
Exceptions	Handles the ones you designed for, stops or escalates on the rest	Handles them, within limits
Determinism	Yes, the same input follows the same predefined logic	No, path can differ
Marginal cost per run	Usually lower on stable, high-volume paths	Usually higher because of model inference; verify against your run volume and exception rate
How it fails	Rule and integration errors, usually reproducible	Decision and tool-use errors, can vary between runs
Best at	High-volume, rule-fixed steps	Judgement calls, dynamic planning, triage where the next step is not fixed
Audit story	Replay the script	Log inputs, outputs, tool calls, approvals, and resulting actions

One row does more work than the rest. Neither kind of automation reliably announces a bad result, so you can’t wait to be told. A script can finish successfully against a wrong rule, and an agent can take a different path on a rerun. What differs is how easily you can reproduce and diagnose the failure afterwards.

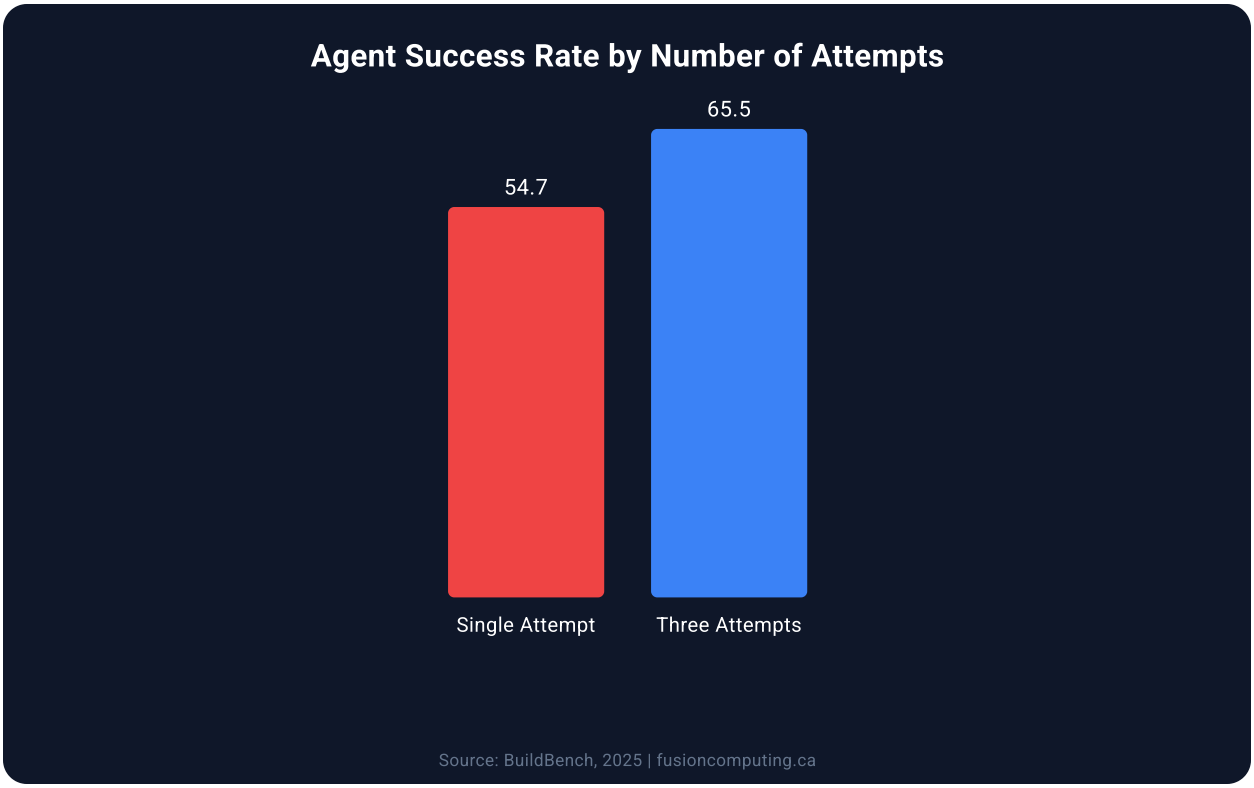

Strict setting, pass@1 versus pass@3. Source: BuildBench, arXiv 2509.25248. Agent Success Rate by Number of Attempts. Source: BuildBench, 2025.

Reproducibility, rather than the vendor category, should decide where each one goes. Put deterministic automation where you need the same result every time and can state the rule exactly. Put agents where the work needs judgement, dynamic planning, or a choice of tool, and a person still checks the output.

What safe agentic automation requires in production

Put the controls in before an agent gets write access, because the alternative is finding out afterwards. Gartner (2026) expects 40% of enterprises to demote or decommission autonomous AI agents by 2027 because of governance gaps identified only after production incidents occur. The order in that sentence is the warning. The gap is found after the incident, which means the controls have to exist before go-live.

Agents are not safe by default, and defaults are what ship unless somebody changes them. Fusion Computing starts with the permissions layer, because it is the part a buyer controls directly.

Gartner's own recommendation in that release is proportional governance, lighter controls for a read-only agent and heavier ones as autonomy rises.

Permissions are the first containment layer, not the whole safety model. They limit what a failure can reach. They do not on their own stop prompt injection, unsafe tool use, or data leaving where it should not, which need guardrails, monitoring, logging, and testing of their own. These are the five controls Fusion Computing puts in first.

1. Give it read access before write access. Most of the value in a first automation is reading, matching, and drafting. It doesn't need write access to do any of that. Let it prepare the work and have a person approve the write. For anything touching money or client records, that approval is the control, not training wheels to remove later.

2. Scope the credentials to the task. An agent doing invoice triage needs the invoice inbox and the AP queue. It doesn't need Global Administrator access in Microsoft Entra. If the only credential available is a broad one, that's what to fix before the automation ships.

> **FIELD NOTE FROM MIKE**
>
> The control that gets dropped first is the approval step. Output looks right for a stretch, someone decides the checking is overhead, and the human review quietly disappears. That review was the control, and it's the one that goes first.
>
> We see the reverse too. A client asks for an agent to do work that is rigid and high-volume, where a fixed flow is cheaper and easier to control.

3. Log every action, and make it reversible where the tool allows. Ask what the rollback is for each tool the agent can call. Where there's genuinely no undo, require approval before execution and write down the recovery path.

4. Run it in parallel before cutting over. Let it work alongside the existing manual process until you have covered enough real cycles and edge cases to judge it. For daily work that is often 2 to 4 weeks. A month-end process needs longer, or replayed historical cases. Quiet failures surface here, and so does an honest accuracy number rather than a vendor's.

5. Set a blast radius. Cap what it can touch per run. An agent that can process 50 invoices can also mis-process 50.

Governance belongs in writing before any of this ships. In Canada that document also has to reflect the privacy law or laws that apply to your organization and data flows. That may include PIPEDA, a provincial private-sector or health-information law, sector-specific rules, or more than one at once.

The rules for what an agent may touch sit alongside the rest of your AI acceptable use policy, and it is worth first checking which AI tools staff are already using without approval, which is what a shadow AI assessment covers.

How to choose what to automate first

Start with the process that has the clearest edges, not the biggest one. In TheAgentCompany benchmark (2025) the most capable agent tested completed 30.3% of 175 tasks inside a simulated software company without help, reaching 39.3% with partial credit. The benchmark's admin and finance categories scored below that average. Treat it as a warning about long-horizon autonomy rather than a failure rate for every office workflow.

> *Verma again, in the same release: "Many use cases positioned as agentic today don't require agentic implementations."*

> *Anushree Verma, Senior Director Analyst, Gartner, Gartner press release (2025)*

Pick the process with the clearest edges, not the biggest one. It's the narrow one that teaches you something.

Good first candidates share three traits. It happens often enough to matter. The inputs are consistent. Someone can tell within a day whether the output was right. Invoice coding, new-starter account provisioning, and moving data between two systems that will never get an integration all qualify.

Want a second opinion before you hand an agent write access? Book a consultation →

> **Why Canadian firms bring this work to Fusion Computing**
>
> CISSP-led, securing IT for Canadian SMBs across Toronto, Hamilton, and Metro Vancouver since 2012.

Bad first candidates are the ones people reach for first. Anything touching payroll. Anything where the rules live in one person's head. Anything where nobody can say what correct means.

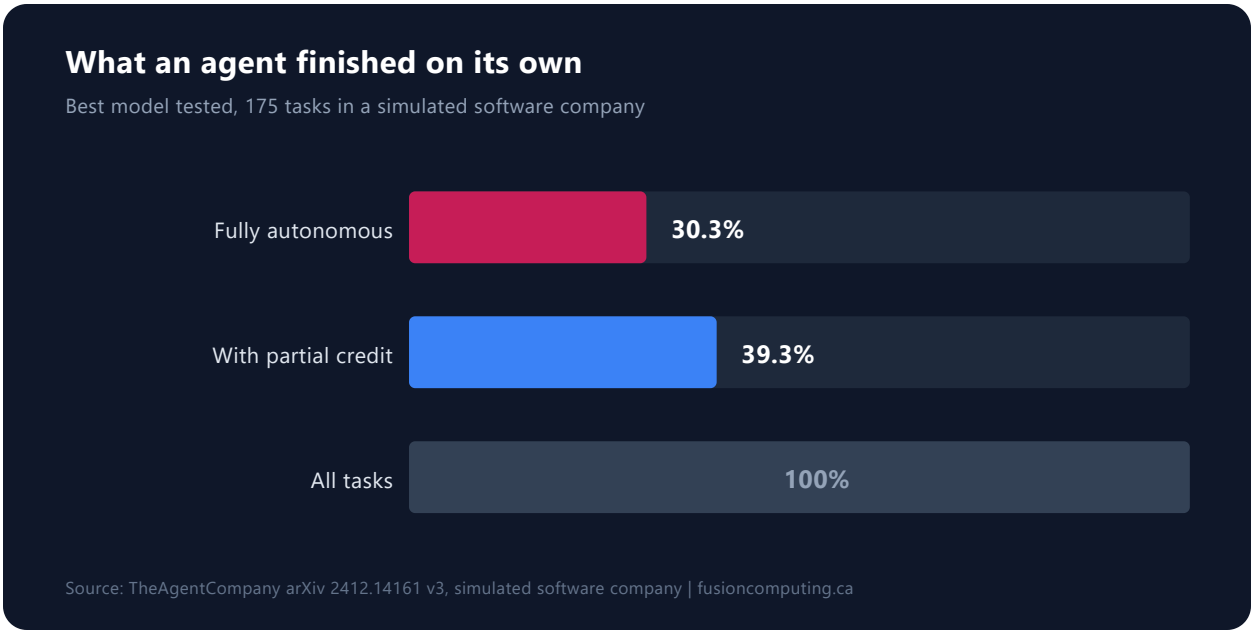

Autonomy on long-horizon tasks in a simulated software company. Source: TheAgentCompany, arXiv 2412.14161 v3.

The upside is real enough to be worth sequencing properly. BDC (2026) reports 30% of Canadian SMEs use generative AI and that those firms describe themselves as 24% more productive, from a survey BDC states is not statistically projectable.

There's an uncomfortable prerequisite. Automation doesn't fix a broken process, it runs it faster. If two people describe the process differently, automating it makes the disagreement arrive sooner. Fusion Computing maps the process before automating it, because in our experience the process a client wants automated usually isn't written down anywhere.

> **What we'd pick:** If Fusion Computing were automating a first process for a 40-person Canadian business today, we'd take the narrowest well-understood task and run a deterministic flow where the rule and path are fixed, or a read-only agent alongside the manual process where the work needs judgement, dynamic planning, or tool selection. One or the other, on the smallest process with clear edges.

This is where the “AI will replace the work anyway” argument breaks down. AI triages, humans own. An agent can run a task well and still leave nobody accountable for noticing when it goes wrong, which is a staffing question, not a software one. The wider sequencing belongs in your AI strategy.

A pre-launch checklist before any automation goes live

Adoption is running ahead of governance in Canadian businesses. Statistics Canada (2026) reports 19.2% used AI to produce goods or deliver services in the 12 months to Q2 2026, triple the 6.1% of 2024, with 13.4% citing cybersecurity or privacy concerns as a barrier. Adoption at that pace is reason enough to put governance in place before an automation goes live.

Pre-launch checklist. Every row needs an answer before an automation touches a production system.

Check	What good looks like
Process written down	One agreed written description of the process, confirmed by the people who do it. If two of them describe it differently, it is not ready
Governance in writing	A written rule for what the automation may touch, reflecting the privacy law or laws that apply to your organization and data flows
Approval gate	Read-only first. Every write to money or client records is approved by a named person before it commits, and that approval stays in place after go-live
Named owner	One person, by name, who receives the alert and knows what the automation was supposed to do
Rollback per tool	Every action is logged, and reversible where the tool allows. Where it is not, approval happens before execution and the recovery path is written down
Parallel-run window	Runs alongside the manual process until enough real cycles and edge cases have been compared, not merely until a date passes
Blast radius	A hard cap on records touched per run, set below what a person could review in a day
Credential scope	Access to the one queue or mailbox the task needs, and nothing wider
Review date	A calendar date to re-check output quality, not an open-ended assumption it still works

The same discipline applies to assistant rollouts, which is why Fusion Computing works through a written Copilot readiness checklist before enabling tenant-wide features.

TRUSTED BY CANADIAN BUSINESSES SINCE 2012

CISSP-led • Securing Canadian SMBs since 2012

What it costs when nobody owns the automation

An unowned agent can produce wrong output quietly while still running up inference cost. Gartner (2026) expects AI inference costs per agentic workflow to rise more than fivefold through 2028, because sophisticated workflows consume far more tokens than simple chatbot exchanges.

The expensive mistake is automating a process nobody owned. That's the one that costs.

When the flow breaks, and it breaks the day a supplier changes an invoice template, somebody has to know it's broken and what it was supposed to do. Automation without an owner becomes invisible infrastructure that fails silently, and nobody's watching it.

> **FIELD NOTE FROM MIKE**
>
> At one of our clients an automation produced wrong output for under a week before anyone noticed. It was caught by a person running a reconciliation, not by the automation reporting a problem.
>
> The tool never signalled a failure at all, because from its point of view nothing had failed. In our experience that is the shape of the risk: the detector is a human on a review cadence, so the control you need is a named owner, not better error handling.

Assigning an owner doesn't need new software, only somebody's time. It's the control we most often find missing. Where there isn't internal capacity to hold it, a co-managed arrangement puts the alert somewhere staffed.

The bottom line

Match the tool to the failure you can afford. If you can't afford a silent one, keep a person on the commit. Rigid, high-volume work wants a script. Work that needs judgement, dynamic planning, or tool selection wants an agent, with a person approving each write. Before either one, write the process down and name who owns the alert.

Talk to Fusion

Related resources

- AI services for Canadian businesses
- Copilot vs ChatGPT vs Claude
- ChatGPT agents and business knowledge
- AI knowledge management playbook
- Microsoft 365 Copilot rollout plan

Frequently asked questions

Q. What is the difference between RPA and an AI agent? –

RPA executes predefined logic. It can branch on rules and handle exceptions you designed for, but it does not improvise outside the paths you built. An AI agent is told what to achieve and chooses the path itself. RPA follows a defined path and is brittle when that path or interface changes. An agent can choose a different path as conditions change, but its outputs are non-deterministic.

Q. Is RPA dead in 2026? –

No. Gartner reports that generative AI and agentic automation slowed RPA market growth in 2024, but the category is still growing. RPA remains the right answer for high-volume, rule-fixed steps between systems that have no API, where determinism and cost per run matter more than flexibility. Paying a reasoning model to re-derive a fixed rule on every run wastes money.

Q. Is it safe to let an AI agent write to our accounting system? –

Not as a first step. Permissions are the first containment layer, not the whole safety model. Run it read-only with a person approving each write, scope its credentials to that task, log every action and make it reversible where the tool allows, cap what it touches per run, and run it in parallel first. Keep that approval after go-live for anything touching money or client records.

Q. Can AI agents work with software that has no API? –

Yes, but that is usually where a deterministic script still wins. For a fixed interface path, RPA avoids paying for model inference on every run and follows a predetermined execution path, so it is often cheaper and easier to audit. Check that against your own run volume and exception rate. Reach for an agent when the work needs judgement, dynamic planning, or tool selection, not merely because the input varies or there is no API.

Q. How do we know if a process is ready to automate? –

If two people in the business describe the process differently, it is not ready. Automation does not fix a broken process, it runs it faster, so a disagreement about how expense approval works simply arrives sooner once automated. Map the process, agree the version, then automate the agreed version. In our experience this mapping step is one buyers often do not budget for.

Q. What should a Canadian small business automate first? –

The process with the clearest edges, not the biggest one. Good first candidates share three traits: it happens often enough to matter, the inputs are consistent, and someone can tell within a day whether the output was right. Invoice coding and new-starter account provisioning are common starting points. Payroll and anything where the rules live in one person’s head are not.

Put this guide to work

Fusion Computing is a CISSP-led Microsoft Solutions Partner and CompTIA Managed Services Trustmark holder, securing Canadian businesses since 2012.

- Managed IT support with a 1-hour priority response
- Cybersecurity, managed detection and response, and compliance reviews
- Microsoft 365 management and licensing

Book a 30-minute scoping call

fusioncomputing.ca/contact-us · 416-566-2845
Toronto · Hamilton · Metro Vancouver