



Professional Services, Software & Solutions Group

GUIDE

The Digital Transformation Guide

A step-by-step playbook: assess, prioritize, modernize, automate and measure — with templates an organization of any size can use this quarter.

PS3G Inc.

1201 N Market Street, Suite 111, Wilmington, DE 19801 · ps3g.com · Email hello@ps3g.com

Contents

- **01** How to use this guide: the eight-step method
- **02** Step 1 – Current-state assessment
- **03** Step 2 – Define outcomes and metrics
- **04** Step 3 – Prioritize initiatives
- **05** Step 4 – Choose the approach per initiative
- **06** Step 5 – Modernize the foundation
- **07** Step 6 – Automate and add AI where it pays
- **08** Step 7 – Deliver in releases
- **09** Step 8 – Measure, govern and keep going
- **10** Transformation readiness checklist

Who this is for

Business owners, founders, department heads and IT leaders who have decided that the way their organization runs today — its systems, its data, its manual work — will not carry it through the next few years. That decision looks the same in a twelve-person firm and a multi-site enterprise; only the scale of the work differs. This guide is written for all of them.

It replaces an earlier edition of the PS3G Digital Transformation Guide. Much of that edition was about moving off paper, adopting cloud applications and building a first integration layer. Those steps still matter, but today the questions have shifted: which processes should be automated, where generative AI and agents pay for themselves, how to keep data fit for those tools, and how to govern all of it under rules that did not exist when the first edition was written.

How to use this document

This is a workbook, not an essay. Each step has a template — a table, a worksheet or a checklist — that you can copy into a spreadsheet or document and fill in with your own information. Work through the steps in order the first time; after that, the method becomes a loop you repeat each quarter.

- **If you have one afternoon:** read Section 01, fill in the systems inventory (Section 02) and the outcomes table (Section 03). You will already know more about your starting point than most organizations do.
- **If you have a quarter:** complete Steps 1 to 4, then plan and ship a first release using Section 08.
- **If you are already mid-way:** use Section 10 to find the gaps and go back to the step that fills them.

A note on examples. Every worked example in this guide is hypothetical and labelled as illustrative. Figures in them are placeholders to show how a template is used, not benchmarks.

01 How to use this guide: the eight-step method

Digital transformation fails most often not for lack of technology but for lack of sequence. Teams buy tools before they know their outcomes, automate processes they have never documented, or add AI on

top of data nobody trusts. The eight steps below put the work in an order that holds up in practice.




Figure 1 – The eight-step method. Steps 1–4 are planning; Steps 5–8 are delivery. The loop is the point.

Three principles that run through every step

Outcomes before tools

No initiative enters the plan without a metric, a baseline and a target. A tool that cannot be tied to one of those is a cost, not an investment.

Foundation before AI

Identity, integration and trustworthy data come first. Every AI feature you add sits on top of them, and inherits their weaknesses.

Releases, not big bangs

Ship something usable every 90 days or sooner. A plan that only pays off in year two will be cancelled in month nine.

Govern as you go

Data ownership, access control and AI oversight are cheaper to design in than to retrofit — and today several of them are required by law or by your customers' contracts.

02 Step 1 – Current-state assessment

You cannot prioritize what you have not written down. The assessment produces three artifacts: a systems inventory, a process pain-point worksheet and a data inventory. Expect it to take a week of part-time effort in a small organization and four to six weeks in a large one with many departments.

Template A – Systems inventory

One row per application, platform or significant spreadsheet. Include the "shadow" tools departments bought on a credit card; they are usually where the real pain lives.

System	Owner	Business function	Users	Hosting	Integrations	Contract / renewal	Health (1-5)
CRM	Sales lead	Pipeline, quotes	14	SaaS	Email only	Annual, March	3
Accounting	Finance	GL, AP, AR	4	SaaS	Bank feed	Monthly	4
Order tracking workbook	Ops manager	Fulfilment status	9	Shared drive	None (manual re-key)	n/a	1
Legacy job system	IT	Scheduling, invoicing	40	On-prem server	Nightly export	Vendor support ended	2

Health score guide: 5 = fit for purpose and supported; 3 = works but with workarounds; 1 = unsupported, fragile or holding the business back. Rows are illustrative.

Template B – Process pain-point worksheet

Interview the people who do the work, not only their managers. Ask each: "What do you re-type, wait for, chase or check by hand every week?" Record the answer in their words.

Process	Where it hurts	Frequency	Hours / week (estimate)	Systems involved	Who is affected
New customer onboarding	Same details entered in CRM, accounting and job system	Daily	6	CRM, accounting, job system	Sales ops, finance
Month-end reporting	Exports merged by hand; numbers disagree	Monthly	20	Accounting, workbook	Finance, leadership
Customer status enquiries	Staff look up three systems to answer one call	Daily	12	All	Customer service

Template C – Data inventory

For each important data set, capture: **name** (customers, orders, contracts, tickets), **system of record** (the one copy that is authoritative), **other copies**, **owner**, **sensitivity** (public / internal / confidential / regulated — for example personal data, health data, payment data), **quality issues** (duplicates, missing fields, stale records) and **retention requirement**. If a data set has no agreed system of record, write "none" — that gap will reappear in Step 5.

Rule of thumb from our engagements. If the inventory turns up more than one system claiming to be the source of truth for customers, fix that before anything else. Nearly every later initiative, and every AI feature, depends on it.

03 Step 2 – Define outcomes and metrics

An outcome is a change in the business that someone outside IT would notice. A metric is how you will know it happened. This step turns "we need to modernize" into a short list of measurable statements that the rest of the plan serves.

Template D – Goal → metric → baseline → target

Keep it to five to eight rows for the whole organization. Each row needs an owner who is accountable for the number, not for the technology. Baselines come from the assessment; if you do not have one, measuring it is the first task.

Area	Goal	Metric	Baseline (today)	Target (12 months)	Owner
Sales	Shorten the quote-to-order cycle	Median days from quote sent to order confirmed	Measure in month 1	Cut by one third	Head of sales
Operations	Eliminate re-keying between systems	Hours per week of manual data entry (from Template B)	Sum of worksheet rows	Reduce by half; remove the order workbook entirely	Operations manager
Finance	Close the books faster and with one set of numbers	Business days to month-end close; number of manual reconciliations	Current close calendar	Close in five days; zero spreadsheet reconciliations	Controller
Customer service	Answer status questions on first contact	Share of enquiries resolved without a call-back; average handling time	Sample two weeks of tickets	Resolve most on first contact	Service lead
Compliance	Be audit-ready without a scramble	Days to produce evidence for a customer security review; open findings	Last review	Evidence pack produced in under a week	Compliance owner
Workforce	Free skilled staff from routine work	Hours per week spent on tasks rated "routine" by the team	From interviews	Redirect to higher-value work	Department heads

Writing a good metric

It is observable without the new system

"Days to close" can be measured today with a calendar. "Adoption of the new platform" cannot — and it measures activity, not outcome.

It has a direction and a date

"Improve reporting" is a wish. "Leadership receives the monthly pack on business day five by Q3" is a target someone can miss.

It is cheap to collect

If the metric requires a project of its own to measure, choose a proxy you can count weekly and accept the approximation.

It is tied to a person

The owner is the department head who benefits. IT owns delivery; the business owns the result. This split matters most in larger organizations where the two are separate teams.

Avoid the vanity row. "Deploy AI across the business" is not an outcome. If an AI initiative belongs in the plan, it will show up as a faster cycle time, lower handling time or fewer errors in one of the rows above — and that is where it should be measured.

04 Step 3 – Prioritize initiatives

The assessment produces a long list of possible initiatives. Most organizations can run two or three at once — a small business perhaps one. Scoring turns a debate about opinions into a conversation about evidence, and it gives you a defensible answer when a department asks why its project is not first.

Template E – Value / feasibility / risk scoring

Score each candidate from 1 (low) to 5 (high) on the three dimensions. **Value** is contribution to the Step 2 outcomes. **Feasibility** covers skills, dependencies, data readiness and cost relative to budget. **Risk** covers disruption, regulatory exposure and the chance of failure. $\text{Priority} = \text{Value} + \text{Feasibility} - \text{Risk}$. Adjust the weights if your situation demands it, but write the weights down and apply them to every row.

Initiative	Outcome served	Value	Feasibility	Risk	Priority	Decision
Integrate CRM and accounting; retire the order workbook	Ops re-keying; finance close	5	4	2	7	Release 1
Single sign-on and MFA across all SaaS tools	Compliance; audit readiness	3	5	1	7	Release 1 (foundation)
Customer status portal with AI assistant	Customer service first-contact resolution	4	2	3	3	Release 3, after integration
Replace legacy job system	Operations; scheduling	5	2	4	3	Plan in Release 2; deliver over two quarters
AI document extraction for supplier invoices	Finance manual entry	3	4	2	5	Release 2
Company-wide chatbot on policy documents	None identified	1	4	2	3	Park until an outcome owns it

Illustrative example for a hypothetical services firm. Scores are placeholders to show the method.

Reading the example

Two things in the table are typical. First, the highest-scoring items are unglamorous: an integration and an identity project. They score well because they are feasible, low-risk and unblock almost everything else. Second, the two most-requested items — the AI assistant and the system replacement — rank low, not because they lack value but because they depend on work that has not been done. The scoring makes the sequence visible: integration first, then the portal that needs integrated data.

Rules for the scoring session

- **Score in one room, in one sitting**, with the outcome owners from Step 2 present. Remote is fine; asynchronous is not — the discussion is the value.
- **Any initiative with no outcome in the second column is parked**. It may be a good idea; it is not yet a priority.
- **Feasibility includes data readiness**. If the data an initiative needs scored poorly in Template C, feasibility cannot be above 3.

- **Re-score every quarter.** Feasibility changes as foundation work lands; risk changes as regulation and contracts change.

05 Step 4 – Choose the approach per initiative

Every prioritized initiative needs a delivery approach, and the right one is rarely "build it ourselves" or "buy the biggest suite". Today there are five realistic options, and most transformations use all of them in different places.

Template F – Approach decision table

Approach	Choose it when	Avoid it when	Typical cost shape	Watch out for
Configure SaaS	The process is standard across your industry and a mature product covers it; your differentiation is elsewhere	The process is how you win; the product forces you to change it	Subscription per user; low upfront	Data lock-in; export rights; per-seat cost growth as you scale
Integrate	The systems are fine individually; the pain is between them (re-keying, mismatched records)	One of the systems is itself end-of-life	One-time build plus modest run cost	Undocumented APIs; who owns the integration when a vendor changes its API
Modernize	A custom or legacy system carries real business logic worth keeping but runs on an unsupported stack or on-prem hardware	Nobody can explain what the system does; a SaaS product now covers it	Project cost, then lower run cost in cloud	Rewriting everything at once; losing edge cases hidden in old code
Build custom	The capability is specific to your business and no product fits without heavy workarounds	A product fits well enough; you cannot commit to maintaining it	Highest upfront; you own the roadmap	Scope creep; building commodity features (login, billing) that products already solve
Add AI	A task involves reading, classifying, drafting or summarizing unstructured content, and a human can check the output	The task needs a guaranteed correct answer and no human review is possible; the data is not ready	Usage-based model costs plus integration	Treating AI as a product instead of a feature; skipping evaluation and guardrails (Section 07)

How the approaches combine




Figure 2 – A typical target shape today: buy the commodity, integrate everything, keep custom work for what differentiates you, and add AI as features on that base.

Decision discipline. Record the approach and the reason for each initiative in one line. When a vendor demo six months later makes "build custom" look unnecessary — or necessary — you will want to know what you were thinking.

06 Step 5 – Modernize the foundation

The foundation is the set of shared capabilities every initiative draws on: where things run, who can access them, how data moves and how it is protected. Foundation work rarely has a department cheering for it, which is why it gets skipped — and why automation and AI projects later stall on access problems, missing data and security findings.

Template G – Foundation baseline checklist

Tick what is already true. Anything unticked that a prioritized initiative depends on goes into Release 1.

Cloud and hosting

- ☐ No business-critical system runs on a server only one person knows how to restart
- ☐ A cloud platform (Azure, AWS or Google Cloud) is chosen for new workloads, with a documented reason
- ☐ Environments (dev, test, production) are separate and reproducible
- ☐ Backups exist, are tested by restoring them, and are stored outside the primary environment
- ☐ Cloud spend is visible monthly by workload

Identity and access

- ☐ One identity provider; single sign-on for every SaaS tool that supports it
- ☐ Multi-factor authentication for all staff, no exceptions for leadership
- ☐ Joiner / mover / leaver process: access removed the day someone leaves
- ☐ Shared passwords and generic admin accounts eliminated
- ☐ Access reviewed each quarter for systems holding sensitive data

Data platform

- ☐ Each core data set (Template C) has one named system of record
- ☐ A central reporting store (warehouse or lakehouse) receives copies from operational systems
- ☐ Key definitions are written down: what counts as a customer, an order, an active account
- ☐ Data quality is measured for the fields the Step 2 metrics depend on
- ☐ Retention and deletion rules exist for regulated and personal data

Integration

- ☐ An integration approach is chosen: integration platform, event bus, or managed API layer
- ☐ No integration depends on a file someone emails or drops on a shared drive
- ☐ Every integration has an owner, monitoring and an alert when it fails
- ☐ APIs for your own systems are documented and versioned

Security baseline

- ☐ A recognized framework is chosen as the yardstick — NIST CSF 2.0 for most organizations; SOC 2 if customers ask for a report; HIPAA where health data is handled — and gaps against it are listed
- ☐ Endpoints, cloud accounts and SaaS admin consoles are logged centrally; someone reviews alerts
- ☐ Vulnerability scanning runs on anything internet-facing; findings have owners and due dates
- ☐ An incident response plan exists on one page and has been walked through once
- ☐ Vendors holding your data are listed with their contracts and security evidence

Right-sizing. A ten-person firm meets most of these with a well-configured identity provider, two or three SaaS products and a managed cloud account. An enterprise meets them with platform teams and tooling. The list is the same; the depth is not. Do not let the enterprise version of a control stop you from doing the simple version today.

07 Step 6 – Automate and add AI where it pays

With the foundation in place, automation and AI stop being experiments and become features of the processes you mapped in Step 1. Work up the ladder below from left to right: each rung is cheaper and more predictable than the next, and most of the value in a typical organization sits on the first three.

Rung	What it is	Good first candidates	Preconditions	Guardrail
1. Workflow automation	Rules-based flows between systems: when X happens, create Y, notify Z	Onboarding, approvals, status notifications, scheduled reports	Integrations exist; the process is documented and stable	Every automation has an owner and a failure alert
2. Document processing	AI extracts structured fields from invoices, forms, contracts, emails	Supplier invoices, intake forms, purchase orders, claims	Sample documents; a place to put the extracted data; a review queue	Human review above a confidence threshold or a value threshold
3. Copilots	Assistants inside the tools staff already use: draft, summarize, suggest	Customer replies, meeting notes, proposal first drafts, code	Clear usage policy; data classification so confidential data stays out of consumer tools	Human sends, approves or commits every output
4. RAG knowledge bases	A model answers questions from your own documents, retrieved at query time from a vector index	Policies and procedures, product documentation, support history	Curated source documents with owners; access control carried into retrieval	Answers cite sources; unanswerable questions say so; evaluation set maintained
5. Agents	A model plans and executes multi-step tasks using tools and APIs	Narrow, reversible tasks: triage tickets, reconcile records, prepare a draft order for approval	Rungs 1–4 working; well-documented APIs; sandbox environment	Least-privilege tool access; spend and action limits; approval gate before irreversible actions; full audit log

Choosing AI candidates: four questions

- 1 Is there an outcome row?** The task must map to a Step 2 metric. If it only maps to "we should be using AI", park it.
- 2 Can a person check the output in less time than doing the task?** If reviewing the result takes as long as producing it, the AI adds cost, not value.
- 3 Is the data ready and allowed?** Clean enough to be useful, classified so you know what the model may see, and permitted by your contracts and by regulations such as HIPAA or the EU AI Act where they apply.
- 4 What happens when it is wrong?** Write the answer down. If the answer is "a customer is harmed" or "money moves", the design needs a human approval step and the initiative moves up a risk point in Template E.

Minimum guardrails for any AI feature

- **An evaluation set** — a few dozen real examples with known good answers, re-run on every change to the prompt, model or data. This is the single practice that separates AI features that keep working from those that quietly degrade.
- **A model-access policy** covering which providers may be used, what data may be sent, and how keys are managed. Use enterprise API agreements that exclude your data from training.
- **Logging** of inputs, outputs and the human decision taken, retained according to your data rules.
- **A named owner** for the feature and a documented way for staff to report a bad output.
- **Prompt-injection and data-leak testing** for anything that reads untrusted content or has tool access — OWASP's Top 10 for LLM Applications is a practical checklist.

08 Step 7 – Deliver in releases

A release is a bundle of changes that reaches real users on a fixed date. Ninety days is long enough to deliver something that matters and short enough that the organization stays interested. Plan the first release in detail, the second in outline, and nothing beyond that in more than a sentence.

Template H – 90-day release plan

Weeks	Phase	Activities	Exit criteria
1–2	Mobilize	Confirm scope from Template E; name the release owner; set up environments, access and the project board; agree the metric readings that define success	One-page release charter signed off by the outcome owner
3–6	Build (first half)	Foundation items first (identity, integration, data fixes); then the initiative's core flow; weekly demo to the outcome owner	Core flow works end-to-end in the test environment with real data
7–10	Build (second half) and pilot	Edge cases, reporting, documentation; pilot with one team or site; collect defects and change requests; train the pilot group	Pilot group can do its job in the new way without the old one
11–12	Roll out	Wider rollout; old process switched off on a named date; hypercare support; first metric reading	Old workbook / process retired; metric reading recorded in the Step 2 table
13	Review and plan next	Retrospective; re-score Template E; draft the next release charter	Next release charter agreed

Roles: small team versus large program

Role	Small organization (one release at a time)	Large program (several releases in parallel)
Sponsor	Owner or managing director; decides scope and spend	Executive sponsor with a steering group meeting monthly
Outcome owner	The department head whose metric moves; often also the pilot lead	One per initiative; accountable for the metric, adoption and the switch-off date
Release owner / delivery lead	An internal lead or your delivery partner's project manager	Program manager plus a delivery lead per workstream
Engineering and integration	A partner's small dedicated team, or one or two internal developers	Platform team for the foundation; product teams per initiative
Data, security and AI governance	Part-time; often the same person as IT lead, supported by an advisor	Named leads for data, security and AI, with a review gate before each release
Change and training	The outcome owner and a "super-user" per team	Change lead; champions network; training built into rollout

Change management checklist

- ☐ Staff who do the work today were interviewed in Step 1 and see their words in the plan
- ☐ The "why" is stated in terms of their day, not the technology
- ☐ A switch-off date for the old process is set before rollout begins and is not negotiable after it
- ☐ Training happens in the week before go-live, on real tasks, not slides
- ☐ A hypercare channel exists for two to four weeks after go-live with same-day answers
- ☐ Wins are reported back to the people who made them happen, with the metric attached

09 Step 8 – Measure, govern and keep going

Transformation does not end; it becomes how the organization runs. The quarterly review is the mechanism. It takes two hours, uses the templates you have already filled in, and produces the next release charter.

Template I – Quarterly review agenda

Item	Input	Output
Outcome scoreboard	Template D with this quarter's readings	Each row marked on track / off track / target reached; off-track rows get a named action
Benefits tracking	Hours freed, processes retired, systems decommissioned, contracts cancelled, findings closed — counted, not estimated	A running benefits ledger the sponsor can read in a minute
Release retrospective	What shipped, what slipped, what broke	Two or three process changes for next release
Re-score the backlog	Template E with updated feasibility and risk	Next release scope
Governance check	Data, security and AI items below	Open risks with owners and dates

Data and AI governance basics

Governance today is not optional for most organizations: customers put it in contracts, and laws such as the EU AI Act and the Colorado AI Act (SB 24-205, now taking effect) create obligations for organizations deploying AI in consequential decisions. The good news is that a proportionate program is small. The elements below scale from a one-page policy in a small firm to a formal management system aligned with ISO/IEC 42001 in a large one.

Data governance

- Owners named for each core data set, with the right to say what it means and who may use it
- Classification scheme applied to systems and shared drives
- Retention and deletion carried out, not just written down
- Quality metrics for the fields that feed decisions and AI features

AI governance

- An inventory of every AI use — including features switched on inside SaaS tools
- A risk tier per use, following NIST AI RMF 1.0: who is affected and how badly if it is wrong
- Human oversight, evaluation and logging proportionate to the tier (Section 07)
- Transparency to staff and customers where the law or common sense requires it

When to bring in outside help

Many organizations run this method entirely with their own people. Outside help earns its place in four situations: when the assessment needs an independent eye because internal views are entrenched; when a release needs engineering capacity the team does not have and will not need permanently; when a specialist area — cloud migration, data platform, AI evaluation, security and compliance — would take longer to learn than the release allows; and when leadership wants an advisor who has seen the pattern before and can say which shortcuts are safe. Whatever the arrangement, keep the outcome ownership, the metrics and the governance decisions inside your organization. A partner can deliver a release; it cannot own your results.

10 Transformation readiness checklist

Print this page. Tick what is true today. Each unticked line points to the step that fills it. An organization that can tick the whole page is not finished — it is ready to run the loop again.

Assess (Step 1)

- ☐ Systems inventory complete, including shadow tools
- ☐ Pain points captured in the words of the people doing the work
- ☐ Every core data set has a named system of record, or the gap is recorded

Outcomes (Step 2)

- ☐ Five to eight outcomes, each with a metric, baseline, target and owner
- ☐ Metrics can be read today, without the new systems
- ☐ No row describes a technology instead of a business result

Prioritize (Step 3)

- ☐ Candidates scored on value, feasibility and risk with written weights
- ☐ Initiatives with no outcome parked, not scheduled
- ☐ Dependencies on foundation work made explicit in the sequence

Approach (Step 4)

- ☐ Each initiative has one approach and a one-line reason
- ☐ Commodity functions are bought, not built
- ☐ Custom work is reserved for what differentiates the business

Foundation (Step 5)

- ☐ Single sign-on and MFA everywhere; leaver access removed same day
- ☐ Backups tested by restoring; no single-person servers
- ☐ Integrations monitored and owned; no emailed-file integrations
- ☐ Security gaps listed against a named framework with owners

Automate and AI (Step 6)

- ☐ Every AI use maps to an outcome row and passes the four questions
- ☐ Evaluation set, logging and human review in place per feature
- ☐ Model-access and data-classification policy published to staff

Deliver (Step 7)

- ☐ A 90-day release charter signed by the outcome owner
- ☐ Roles filled for sponsor, outcome owner, delivery, governance, change
- ☐ Switch-off date for the old process set before rollout

Measure and govern (Step 8)

- ☐ Quarterly review scheduled with the scoreboard and benefits ledger
- ☐ AI inventory and risk tiers maintained; data owners named
- ☐ Next release charter drafted before the current one closes

How to score yourself. Fewer than a third ticked: start at Step 1 and give the assessment the time it needs. Between a third and two-thirds: you have a foundation; pick the weakest step and make it the theme of the next release. More than two-thirds: your constraint is now delivery pace and governance depth — focus Steps 7 and 8.



Professional Services, Software & Solutions Group

About PS3G

PS3G Inc. is a professional services, software and solutions company headquartered in Wilmington, Delaware, with a delivery center in India that lets us work with clients around the clock. We work with organizations of every size — from owner-led businesses to enterprise IT groups — on websites and web applications, mobile apps, .NET and Microsoft engineering, cloud migration and DevOps, data platforms and data governance, AI consulting and development, AI governance, cybersecurity and compliance, software testing, managed services, and IT staffing with dedicated teams.

We also build and run our own products. DocVilla, our HIPAA-compliant cloud EHR, telehealth and practice-management platform, is used by medical practices across the United States — which means the methods in this guide are ones we apply to systems we operate ourselves, not only ones we advise on.

How we engage

Every engagement starts the same way: email us → we reply with questions → a 30-minute video call we schedule → a short assessment → a written proposal with scope, timeline and cost. From there we work in whichever model fits — a fixed-scope project, a dedicated team working alongside yours, managed services for what you would rather not run, or advisory support for the leaders running the program.

Talk to us

If this guide has given you a list of unticked boxes, we can help you work through them — whether that is a current-state assessment, a foundation release, an AI feature built with proper guardrails, or a governance program that will satisfy your customers and regulators.

Email hello@ps3g.com.

PS3G Inc., 1201 N Market Street, Suite 111, Wilmington, DE 19801 · ps3g.com

This guide was written by the PS3G team and supersedes an earlier edition of the PS3G Digital Transformation Guide. It is general guidance, not legal, regulatory or financial advice. Frameworks and laws referenced — NIST AI RMF 1.0, NIST CSF 2.0, ISO/IEC 42001, the EU AI Act, Colorado SB 24-205, HIPAA, SOC 2 and the OWASP Top 10 for LLM Applications — should be consulted directly for current requirements.