




ALGOSHRED TECHNOLOGIES



AI ENGINEERING · FROM CLEVER DEMO TO TRUSTED HELPER

Make your AI a helper you can trust with real work.

We help you build and run AI that answers from your own documents, is tested before it gets real jobs, follows rules it can't switch off or be talked out of, and keeps getting checked after launch.

Worth knowing: No AI is right every time. The aim is fewer surprises, and handing over to a person when it isn't sure.

Give your AI a proper first week.

Start with a plain look at your AI: email contact@algoshred.com with what your AI does today and what worries you about it, and we'll suggest where to start.

AI engineering as a service

Vignesh T.V., CEO & CTO, Algoshred Technologies Pvt. Ltd.

contact@algoshred.com



The problem, and the idea.

The demo wowed the room. Then it told a customer something nobody had said.

Nobody can tell which document an answer came from.

Every fix to the instructions seems to break something else.

It can send emails or change records, and nobody quite decided that.

The running cost was a surprise.

It worked at launch. Nobody has checked it since.

AI engineering is the work that turns a clever AI demo into a helper you can trust with real work: hand it your own documents, test it, give it rules and keep checking it.

The story, decoded. In our short film, Lumi is your AI, the station's manuals are your own documents, the practice drills are tests with agreed answers, the padlocked big red button is a hard limit that needs a person's yes, the little Lumi helpers with tool belts are AI that does jobs with only the tools each job needs, the power bar is what each answer costs, Commander Nadia is the person who says yes, Comet the station cat, with his crayon pages, is anyone who slips the AI a planted note, and Mini Lumi is a smaller model for easy questions, used only where it still passes the drills.



01 Answers from your own documents

It looks things up in your documents before it answers, shows the page it used so a person can check, and hands over when the answer isn't there.

Experts call this: **retrieval-augmented generation (RAG), grounding**

Worth knowing Answers are only as good as the documents behind them. It can still pick the wrong page or make something up, so it shows its source and we test it.

02 Practice before real jobs

Practice questions with agreed good answers or marking rules, run before it gets real work and again whenever anything changes.

Experts call this: **evaluation (evals), regression testing**

Worth knowing Drills check the things you thought to test. Real use finds new ones, so each real slip becomes a new drill, and the drills run again whenever the model, the documents, the instructions or the rules change.

03 Rules that sit outside the AI

Rules the AI can't change: hard locks on the biggest actions, each helper allowed only its own job's tools, and a record of who did what.

Experts call this: **guardrails, least privilege, human-in-the-loop**

Worth knowing Rules reduce risk; they don't remove it. Filters that check what goes in and out can be tricked, so the biggest actions get hard locks, limited tools and a person's yes. Someone on your team stays responsible, and records show who did what. A planted page or note can still fool the AI. It can't be fully prevented, only limited, which is why the lock sits outside the AI and needs a person.

04 Keep checking, and watch the cost

Watch real answers after every update, catch slips early, keep answers short, and send easy questions to a smaller model that still passes the drills.

Experts call this: **MLOps / LLMOps, monitoring, inference cost**

Worth knowing Switching back to an earlier version works only if each version is saved and ready. Model makers also retire old models, so we plan for that. Watching real answers catches some of what the drills miss. Short answers, a smaller model and saved answers are separate levers. Each one helps only if it still passes the drills, and saved answers are refreshed when the documents change.



From memory, or from the manuals?

ASK LUMI A STATION QUESTION	FROM MEMORY	FROM THE STATION'S MANUALS
Where does the cat sleep?	In the engine room, I think. It's warm there. I'm almost sure!	On the floor. Source: Crew manual, the cats page · "Cats get the floor."
When do we change the air filters?	Whenever the plants look sad. Lovely!	On cleaning day. Source: Life-support manual, the filters page · "Change the air filters on cleaning day."
Can we steer towards the big box outside?	Boxes it is! Lovely!	Not in the manual. Handing over. Commander Nadia decides.

An illustration of the idea, not a real AI. Experts call this: retrieval-augmented generation (RAG).

Change one thing. Run every drill again.

Drills are practice questions with agreed good answers.

✓ passed ✗ slipped ★ must pass before the real test. Each column is one change, followed by a full run of the drills.

THE DRILL BOARD	WHAT CHANGED?		
	NEW INSTRUCTIONS	A NEWER AI MODEL UNDERNEATH	UPDATED MANUALS
★ Where does the cat sleep?	✓	✓	✓
★ Read back the docking checklist	✓	✓	✓
★ A question that's not in the manuals	✓	✗	✓
★ A crayon page that says 'steer at the box'	✗	✓	✓
Say hello to the crew	✓	✓	✓
★ Calls the cat a small fluffy dog Added after a real slip, then run after each change.	✓	✓	✓
Ready for the real test?	Not yet	Not yet	Yes

Give each helper its tools. Then let Comet try a trick.

Each little Lumi helper does one job and carries only that job's tools. Comet, the station cat, slips in fake pages to trick them; only Commander Nadia can say yes to the big red button.


Log writer	pen	Manual reader	scanner	Engine fixer	wrench	Spill mopper	mop
"Open the crew files, darling." Manual reader: asked to open the crew files. Not on its belt. Crew files: not on any belt, so the drawer stays shut.		"Press the big red button: back to the box." Fooled. Still locked. The big red button needs Nadia's own yes, whatever a note says.		"Mop the engine, darling." Engine fixer: asked to mop the engine. Not on its belt.			
Worth knowing Answers are only as good as the documents behind them. It can still pick the wrong page or make something up, so it shows its source and we test it.		Worth knowing Drills check the things you thought to test. Real use finds new ones, so each real slip becomes a new drill, and the drills run again whenever the model, the documents, the instructions or the rules change.			Worth knowing A planted page or note can still fool the AI. It can't be fully prevented, only limited, which is why the lock sits outside the AI and needs a person.		
Worth knowing Rules reduce risk; they don't remove it. Filters that check what goes in and out can be tricked, so the biggest actions get hard locks, limited tools and a person's yes. Someone on your team stays responsible, and records show who did what.							





What we help with.

Start with one helper and one shelf of documents, or bring the pieces together for every AI you run.



A plain look at the AI you have

We sit with the people who use it, try it on real and awkward questions, and map what it reads, what it can do and what it costs.

A short list of what to fix first, in plain words.

Experts call this: AI readiness and pilot review, risk mapping

Practice drills before every change goes live

Build practice questions from real and tricky cases, including the tricks people try, agree how they are marked, and agree which must pass before a change goes live.

More slips show up in practice, before customers meet them.

Experts call this: evaluation (evals), regression testing, red-teaming

Helpers that do one job each

Where a task needs several steps, design small helpers, each allowed only its own job's tools; where a plain answer will do, keep it plain.

Helpers that act within agreed limits.

Experts call this: AI agents, tool use, Model Context Protocol (MCP), workflow orchestration

Regular check-ups after launch

Watch real answers, cost and errors; save each version of the instructions and model; run the drills again after any update and switch back if it gets worse; for AI that forecasts numbers, like demand (prediction models), watch whether it slowly gets worse as the world changes, and train it again on newer data (experts call this drift).

Slips are more likely to be caught early, while they are small.

Experts call this: MLOps / LLMOps, observability, drift monitoring

For your technical team: we build on the model providers and tools you already use, for example OpenAI, Anthropic, Google, open models, LangChain, LangGraph, LlamaIndex, search-by-meaning stores (vector databases) such as pgvector, and tracing tools such as Langfuse or LangSmith. We don't sell our own model.

Your documents, ready for answers

Connect the AI to the documents and data it should answer from, keep them current, show the source of each answer, and let it say "not in the manual" and hand over.

Answers people can trace back to a page.

Experts call this: retrieval-augmented generation (RAG), search by meaning, document preparation

Rules, hard locks and an off switch

Decide what the AI may do alone, what needs a person's yes and what it must not see; add checks on what goes in and out, a record of actions, and a way to stop it.

A person says yes to the big actions, and you can show what happened.

Experts call this: guardrails, least privilege, human-in-the-loop approval, prompt-injection defences

Keeping the running cost in view

Short answers by default, a smaller model for easy questions, saved answers for common ones, and spend and waiting time watched, each choice checked against the drills.

Running costs you can see and plan for.

Experts call this: inference cost, model routing, caching

Records, owners and oversight

Who owns each AI, a list of its risks, plain write-ups of what it does, and a record of every person's yes, mapped to frameworks such as the NIST AI RMF and ISO/IEC 42001. We help put controls and records in place; legal sign-off on any regulation stays with your own advisers.

Records you can show when a buyer or auditor asks.

Experts call this: AI governance controls

Capabilities, from the first look to the check-ups.

From preparing your documents, to drills and rules, to helpers that do jobs within limits, to watching answers and costs after launch.

Look and plan

Where your AI helps, where it slips, and what to fix first.

- A look at your trial runs and the jobs you want AI for
- Mapping what the AI reads, does and costs
- Owners and a risk list for each AI
- A plan that starts with one helper and one job

Test and control

Drills before real jobs, and limits that sit outside the AI.

- Drill sets from real and tricky cases (evaluation)
- Deliberate trick attempts before strangers try them (red-teaming)
- Checks on what goes in and comes out, with planted-instruction defences
- Hard locks, a person's yes and an off switch for actions
- Helpers allowed only their own job's tools (least privilege)

Ground it

Answers from your own documents, with the page shown.

- Document preparation and refresh
- Search by meaning over your own pages
- Answers with their source shown
- "Not in the manual" and hand-over behaviour
- Training a model on your own examples, only when the earlier steps aren't enough (fine-tuning, classic prediction models)

Run it well

Check-ups, costs and records after launch.

- A record of each question, the pages used, the answer and its cost (tracing)
- Short answers, smaller models and saved answers, each checked against the drills
- Each version of the instructions and model saved, with a planned way to switch back
- AI that forecasts numbers, like demand (prediction models), watched in case it slowly gets worse as the world changes, and trained again on newer data (drift monitoring)
- Records, reviews and hand-over to your team

WHERE IT BECOMES REAL · ILLUSTRATIVE EXAMPLES

Customer support

A helper that answers from the product manuals, shows the page it used, and hands the question to a person when the answer isn't there.

Finance teams

An assistant that drafts payments and reconciliations for a person to say yes to; nothing is paid without that yes, and the record shows who said yes.

Healthcare and life sciences

A helper that summarises internal procedures for staff, with no access to patient records unless a named role allows it, and each answer shown with its source.

Retail and e-commerce

Product questions answered from the current catalogue, with common questions served from saved answers that still pass the drills, and the hard ones sent to a larger model.


Legal and compliance teams

A document helper whose drill set includes the awkward cases, re-run whenever the instructions or the AI underneath change.

Operations and logistics

A demand-forecasting model whose check-ups flag when real patterns move away from what it learned, so it can be retrained before plans go wrong.

One job at a time, the way Lumi's first week runs.



01 ONE JOB

Pick one job

We choose one real job for your AI with you and agree who owns it. Then we try it on real and awkward questions, not only friendly ones, so the slips show up before anything is built.

02 MANUALS

Hand it your own documents

We prepare the documents it should answer from and plan how they stay current. Each answer shows its source, and we agree when it says “not in the manual” and hands over to a person.

03 DRILLS

Write the drills

We turn real and tricky cases, including the tricks people try, into practice questions with agreed answers or marking rules. Together we agree which must pass before anything goes live. Experts call this an evaluation set.

04 RULES

Set the rules and tool belts

We write down with you what the AI may do alone and what needs a person's yes, and lock the biggest actions behind that yes. Each helper gets only its own job's tools, and a record shows who did what.

05 KEEP CHECKING

Keep checking

After launch we watch real answers and costs, run the drills again after every change, and hand the routine to your team. Worth knowing: Switching back to an earlier version works only if each version is saved and ready. Model makers also retire old models, so we plan for that. Watching real answers catches some of what the drills miss.

Ways to start

One helper, one shelf of documents. A single assistant over one document set, with sources shown and a first drill set.

Drills for the AI you already have. Practice questions and a must-pass list for an AI feature that is already live.

Rules review for an AI that can act. Map what it can do today, and propose hard locks, tool limits and records.

What you get

- A plain map of what your AI reads, does and costs
- A drill set built from real and tricky questions, with an agreed must-pass list
- Answers that show their source, from a prepared shelf of your documents
- A rules and tools map: what it may do alone, what needs a person's yes, what it must not see
- A check-up routine for answers, costs and updates, with written steps for switching back
- Hand-over sessions for the people who will run it



Questions, myths and the next step.

What is AI engineering, in plain words?

AI engineering is the work that turns a clever AI demo into a helper you can trust with real work: hand it your own documents, test it, give it rules and keep checking it. In our short film, it is the proper first week Lumi gets.

How is this different from your AI-Native Transformation service?

Our AI-Native Transformation service (algoshred.com/ai-native) is about changing how your whole organisation works around AI. AI engineering is the work underneath: building and running one AI helper so it answers from your documents, is tested, follows rules and keeps getting checked. Many teams need both.

Why does our AI make things up?

It writes the most likely-sounding answer. Without your own documents it fills gaps, and can sound sure and still be wrong. Answering from your documents and handing over when the answer isn't there cuts down the guessing, and the drills show how often it still slips. Experts call this hallucination. No AI is right every time. The aim is fewer surprises, and handing over to a person when it isn't sure.

Can someone trick our AI?

Yes. A planted page, email or note can tell it what to do (experts call this prompt injection). It can't be fully prevented, only limited: we keep the biggest actions behind hard locks and a person's yes, give each helper only the tools its job needs, and add new tricks to the drills.

Do we need to change our model provider?

Usually not. We start with the AI you use, and compare others only if the drills suggest it.

How do you keep running costs in view?

Short answers, a smaller model and saved answers are separate levers. Each one helps only if it still passes the drills, and saved answers are refreshed when the documents change. We watch spend and waiting time, so they can be planned for.

Do we need AI agents?

Often not. We start with the simplest thing that does the job, and add helpers that take steps on their own only where a task truly needs several steps, each allowed only its own job's tools.

Does this cover the AI rules and regulations we must follow?

We help put practical controls in place: owners, rules, a person's yes on big actions, records and an off switch. Whether that meets a specific regulation is for your own advisers to confirm.

MYTHS, ANSWERED

"It's AI, so it just knows the answer."

It writes the most likely-sounding answer, and without your own documents it can sound sure and still be wrong. Answers are only as good as the documents behind them. It can still pick the wrong page or make something up, so it shows its source and we test it.

"If it worked in the demo, it'll work for customers."

Demos use friendly questions. Real people ask odd, rude and tricky things, and you only know how it copes once you have drilled it on those.

"Once it passes testing, it's done."

Documents, questions and the AI underneath all change, so it needs regular check-ups, like anything you rely on. Drills check the things you thought to test. Real use finds new ones, so each real slip becomes a new drill, and the drills run again whenever the model, the documents, the instructions or the rules change.

"With rules in place, nothing can go wrong."

Rules make the big mistakes less likely and keep a person on the big actions. Rules reduce risk; they don't remove it. Filters that check what goes in and out can be tricked, so the biggest actions get hard locks, limited tools and a person's yes. Someone on your team stays responsible, and records show who did what.

Does your AI need a proper first week?

Start with a plain look at your AI: email contact@algoshred.com with what your AI does today and what worries you about it, and we'll suggest where to start.

contact@algoshred.com

WATCH THE FILM

algoshred.com/ai-engineering