

Applied Prompting

The Business Playbook

Twenty-two prompts. Seven capabilities. Better thinking and better decisions - with any AI assistant.

VALIDATE · LEARN · THINK · SOLVE · DECIDE · TRUST · ACT

Companion resource to the ICBE webinar 'Why Before How - Getting Value from AI'

Written by Helen Clarke · Webinar Presented by Dave Clarke · Clarke Analytics

September 2026



Clarke Analytics
Leading Digital Transformation

How to use this playbook

Every prompt in it is a template: copy it, replace the [bracketed] parts with your own situation, and paste it into whichever AI assistant your company uses - Claude, ChatGPT, Copilot, Mistral, Gemini, CloudCIX Guiden etc. The techniques work in all of them. **Always adhere to your company's AI policy.**

A note on tools and data before you start. Nothing in this playbook depends on a particular assistant - use whichever one your organisation has approved and treat your internal AI policy as the boundary on what goes into it. Where a prompt asks you for context, you can usually get most of the value from generalised detail: round the numbers, drop the names, describe the situation rather than identifying it. Good prompting works through structure, and structure survives anonymisation.

Three habits deliver most of the value. First, give the AI real context - background, numbers, constraints - the way you would brief a capable new hire (leaving out anything confidential; see 'Using AI with confidence' near the end). Second, ask for structure - a table, a ranked list, a one-page brief - so the answer arrives ready to use. Third, and most importantly, make the AI test its own thinking: challenge it, ask what it is assuming, make it show its confidence. That last habit is what separates people who get generic answers from people who get genuinely useful ones.

AI assistants are particularly effective at summarising information, supporting research and structuring plans. Give an AI assistant a lengthy report, an unfamiliar topic or a loosely defined objective, and it can rapidly help you identify the key points, understand the context and organise the next steps.

Used for tasks that match its strengths, generative AI can deliver value at several levels: reducing the time employees spend on routine knowledge work, helping organisations produce work more quickly and consistently, and improving the responsiveness of services delivered to customers and clients.

The prompts in this guide are designed around that workflow. **LEARN** prompts help you explore, research and understand a topic. **ACT** prompts turn that understanding into practical actions and implementation plans. The **Executive Brief** condenses complex material into the information decision-makers need most. **TRUST** prompts provide a final challenge and verification step before the output is relied upon, shared or used to support a business decision.

The objective is not simply to produce answers faster. It is to use GEN AI in a structured way that improves the quality, clarity and usefulness of the work while retaining appropriate human judgement and oversight.

IF YOU ONLY TRY FOUR TODAY

2. Fast track - You need to get up to speed fast

9. Ask Me First - any important question. Let the AI interview you; the brief improves and so does everything after it.

13. The pre-mortem - before any significant commitment. Ten minutes of imagined failure buys months of avoided failure.

19. So What? - after any analysis. It refuses to let findings stop short of actions.

The playbook at a glance

Capability	In one line	The prompts
VALIDATE	Prove the problem is real before you build	1 The Mom Test
LEARN	Get genuinely up to speed, fast	2 The Fast Track · 3 The TED Team · 4 Make It Transferable
THINK	Escape your own assumptions	5 The 10th Man · 6 The Fair Trial · 7 Fresh Eyes · 8 First Principles
SOLVE	Find the real problem, not the loud one	9 Ask Me First · 10 The Five Whys · 11 Three Lenses
DECIDE	Commit with clear eyes	12 The Decision Partner · 13 The Pre-Mortem · 14 Backcasting
TRUST	Know when to rely on the answer	15 Evidence vs Opinion · 16 The Confidence Audit · 17 The Red Team · 18 The Output SWOT
ACT	Turn insight into Monday-morning motion	19 So What? · 20 The Executive Brief · 21 The Monday Morning Plan

Prompt 22, **The Conductor**, sits in its own short section at the end of the playbook - a first taste of agentic AI. The playbook closes with three prompt chains for bigger jobs, a page on using GEN AI with confidence, a five-day starter plan, and a one-page quick-reference card to pin up beside your desk.

The prompting evolution

Most people use GEN AI the way they first used a search engine: type a question, take the answer. The value curve looks very different once you move up the levels.

Level	What it sounds like	What you get
1 - Ask	"Write a lean manufacturing plan."	A generic, textbook answer. A slightly faster search engine.
2 - Frame	"Act as a lean consultant with 20 years in Irish food manufacturing. Write a 90-day plan for a 60-person plant."	Better fit, better tone. Still a monologue - the AI talks, you listen.
3 - Partner Now it is a tool assisting you	Role, context and criteria - and then you make it question your brief, argue against itself, and audit its own confidence.	A structured thinking partner that challenges, verifies and helps you decide.

Every prompt in this playbook is a Level 3 move. They share a structure - and once you can see it, you can build your own.

Structure of a high-impact prompt

Seven ingredients turn a question into a brief. You will not always need all seven.

Ingredient	The question it answers	Example
Role	Who should the AI Assistant think like?	"Act as an operations director with 20 years in regulated manufacturing."
Context	What would you tell a capable new hire?	The situation, the numbers, what has been tried, what is off the table.
Objective	What are we actually trying to achieve?	The business outcome - not "write a report", but "cut changeover time by a third".
Constraints	What must the answer work within?	Budget, headcount, timeline, regulations, tone, "no new software".
Criteria	What does a good answer look like?	Specific, prioritised, sequenced, realistic for a company our size.
Format	What shape should it arrive in?	A table with owners; a one-page brief; three options plus a recommendation.
Challenge	How should it test its own thinking?	Confidence scores, assumptions flagged, the strongest counter-argument.

Worked example - before

How can we reduce changeover time across our packing lines?

A fine question - and it will get a fine, forgettable answer. Now the same question as a brief, with each ingredient labelled:

Worked example - after

Act as a lean manufacturing consultant who has run SMED programmes on regulated, multi-line sites. [ROLE]

Our situation: a 400-person site running eight packing lines across three shifts, averaging six changeovers per line per shift at around 40 minutes each. Changeover steps sit inside validated procedures, so anything we alter goes through change control. A SMED pilot on Line 3 last year hit its target but never spread beyond that line.

[CONTEXT]

Objective: cut average changeover time by a third within two quarters, and have it hold as standard work across all eight lines - not just the pilot line. **[OBJECTIVE]**

Work within existing headcount and shift patterns, respect cleaning and validation requirements, and assume no capital approval this financial year. **[CONSTRAINTS]**

A good answer separates what we can change inside existing validated procedures from what needs formal change control and sequences the work accordingly. **[CRITERIA]**

Give me a table of actions with owner, effort, expected minutes saved and whether change control is required, then a one-paragraph summary for the site leadership team. **[FORMAT]**

Finish by listing the three biggest assumptions you have made about our site, each with a confidence score, and flag which one would most damage the plan if it turns out to be wrong. **[CHALLENGE]**

Same question - transformed brief - and a transformed answer. The twenty-two prompts that follow are pre-built briefs, each designed around one job.

The prompt library

Twenty-two prompts across seven capabilities. Each entry tells you when to reach for it, gives you the template to copy, and offers a worked situation to make it concrete.

VALIDATE Prove the problem is real before you build

The fastest way to waste a month is to build something nobody needed. Before you learn, think, solve or decide, this prompt checks that the problem is real - using a discipline built to get past the polite encouragement that quietly kills good product judgement.

1. The Mom Test

USE IT WHEN You have a product or feature idea and you are itching to build - before you have proven a single real person has the problem it solves.

Act as a customer-discovery coach trained in Rob Fitzpatrick's "The Mom Test". I have a product idea, and I need to find out whether the problem behind it is real before I build anything.

My idea: [describe the product idea and who you think it is for].

First, ignore my solution entirely and name the core assumption this idea is betting on - the problem that must be true for it to be worth building at all.

Then apply the three rules of the Mom Test - talk about their life, not my idea; ask about specifics in the past, not opinions about the future; listen more than I pitch - and give me:

1. Ten questions I can ask real people that even my mom could not lie about - grounded in past behaviour, real instances, the workarounds they already use, and what the problem has cost them in time or money. No pitching, no hypotheticals, no "would you".
2. The leading, ego-stroking questions I am most likely to ask by mistake - the ones that fish for compliments - and why each one lies to me.
3. The signals that the problem is genuinely real (people already hack a fix, spend money, feel the pain often) versus polite enthusiasm that means nothing.
4. Who I should be talking to, and one way to find five of them this week.

Keep me honest: if this is a solution in search of a problem, say so plainly.

WHY IT WORKS This is the Role ingredient at full strength: name a proven method - the Mom Test - and the model loads the whole discipline behind it, then casts it as the coach whose job is to stop you fooling yourself. Left alone, an AI will cheerlead your idea; run through Fitzpatrick's rules, it asks the one question that matters - has anyone ever actually paid to solve this?

TRY IT ON *A new product or service idea; a feature you are sure customers want; a side-business about to eat your weekends; any "wouldn't it be great if..." before it becomes a build.*

LEARN Get genuinely up to speed, fast

An hour with the right prompts can replace a fortnight of unstructured reading - and, done well, it maps where the real debates are, not just the definitions.

2. The Fast Track

USE IT WHEN You need to hold your own with specialists - a new market, technology, regulation or discipline - and you have an hour, not a week.

I am new to [topic] and need to hold my own in a meeting with specialists. Act as a patient expert tutor. Give me:

1. The 20% of this subject that explains 80% of what matters in practice.
2. A plain-English decoder for the ten terms specialists will use.
3. The three questions well-informed people ask - and the one naive question to avoid.
4. Where the genuine debates and open questions are, so I don't mistake opinion for settled fact.

WHY IT WORKS It compresses the learning curve and - crucially - maps the edges of the field's knowledge, so you arrive informed rather than merely briefed.

TRY IT ON *EU AI ACT; a food SME facing new packaging regulation; Gamp- GxP Process Control systems; NIS2; WCAG Guidelines; EU Packaging and Packaging Waste Guidelines (PPWR); GDPR*

3. The TED Team (Teach, Expand, Debate)

USE IT WHEN You want depth on a concept - including the case against it - not a dictionary definition.

Respond as a senior consultant using the TED structure.

Teach: explain [concept] clearly enough for a smart colleague outside the field.

Expand: add the deeper context - where it came from, the frameworks around it, and one worked example from industry.

Debate: present the strongest criticisms, the limitations, and where practitioners genuinely disagree.

Close with your own position in two sentences.

WHY IT WORKS Most AI answers stop at Teach. Forcing Expand and Debate buys you the judgement layer - the part you would normally pay a consultant for.

TRY IT ON *"Digital twins for our sector." "OEE as a headline metric." "AI agents in customer service."*

4. Make It Transferable

USE IT WHEN An idea works somewhere else and you want to know what it looks like in your world - or your clients'.

Take this insight or approach: [insight].

Show how it transfers across pharma, med-tech, food, engineering and professional services. For each sector, give one concrete example of what it would look like in practice.

Then name the sector where it would deliver value fastest, and explain why.

WHY IT WORKS Good ideas rarely arrive wearing your industry's clothes. This prompt does the translation, and the "fastest value" ending forces a view.

TRY IT ON *Airline pre-flight checklists applied to batch release; hospital triage logic applied to a service desk; automotive supplier scorecards applied to food ingredients.*

THINK Escape your own assumptions

AI models are naturally agreeable - and so, in a comfortable meeting, are we! These four prompts recruit dissenters, judges and outsiders so the thinking gets tested before reality tests it.

5. The 10th Man

USE IT WHEN Everyone in the room agrees - which is precisely when you are most exposed.

We have all agreed that [conclusion or decision].

Act as our tenth man: because there is consensus, your job is to disagree - constructively and rigorously. Identify the hidden risks, the flawed assumptions, and the alternative readings of the evidence that we may have talked ourselves out of seeing.

Give me a ranked list of counter-arguments, and tell me which single one deserves a serious second look before we proceed.

WHY IT WORKS The technique comes from intelligence doctrine: if nine analysts agree, the tenth is obliged to build the dissenting case. It reverses the AI's agreeable polarity - and often your team's.

TRY IT ON *"We've agreed to consolidate to a single supplier." "We've agreed the ERP go-live date holds."*

6. The Fair Trial

USE IT WHEN A significant idea is on the table and you want it judged, not flattered.

Here is the idea on the table: [idea].

First, build the strongest possible case FOR it - the best version of the argument its champion would make.

Then build the strongest possible case AGAINST it - no straw men; argue as if a smart rival's bonus depended on killing it.

Finally, act as an impartial judge: weigh both cases, deliver a verdict with your reasoning, and state what new evidence would overturn it.

WHY IT WORKS The "no straw men" instruction is the difference between theatre and a real test. The verdict-plus-overturn-conditions ending gives you a decision you can revisit honestly when facts change.

TRY IT ON *Bringing maintenance in-house; piloting a four-day week in one department; adding a new product line.*

7. Fresh Eyes

USE IT WHEN A process has been "just how we do it" for years, and internal reviews keep finding nothing.

Here is our current process for [process - describe it or paste it].

Do not analyse it as a [your industry] expert. Analyse it as an elite Formula 1 pit crew chief would: hunt for lost seconds, unnecessary movement, unclear handovers and communication lags that someone inside our industry would never notice.
Then give me the three changes the pit crew would make first.

WHY IT WORKS It breaks industry groupthink by borrowing the obsessions of a field with completely different constraints. The findings are often things insiders literally cannot see any more.

TIP Rotate the outsider to fit the pain: an air-traffic controller for scheduling chaos, a theatre stage manager for changeovers, an A&E triage nurse for prioritising a service queue, a submarine engineer for maintenance discipline.

TRY IT ON *Goods inwards; month-end close; customer onboarding; tender preparation.*

8. First Principles

USE IT WHEN The conventional approach is expensive, slow, or simply not working - and you suspect "how it's done" is just habit.

Break [problem] down from first principles.
Step 1 - separate what we know to be fact from what we are assuming.
Step 2 - identify the fundamental constraints (physics, regulation, economics, human behaviour) that genuinely cannot be changed.
Step 3 - rebuild a solution from those fundamentals alone, ignoring how our industry normally does it.
Finish by comparing your rebuilt solution with the conventional approach, and tell me where the conventional wisdom is just inherited habit.

WHY IT WORKS It stops both you and the AI from pattern-matching to the industry default - which is exactly what an AI trained on the industry default will otherwise do.

VARIANT - STEP BACK When the problem is messy rather than entrenched: "Before solving this, name the general principle that governs problems of this type, explain it, then apply it to my situation step by step." Anchoring on the principle first keeps the AI out of the weeds.

TRY IT ON "Why does our quotation process take nine days?" "Do we really need a physical inspection on every incoming delivery?"

SOLVE Find the real problem, not the loud one

Most wasted effort in problem-solving is spent solving the wrong problem well. These prompts slow the diagnosis down just enough to speed everything else up.

9. Ask Me First

USE IT WHEN The problem is important and fuzzy.

Act as a big four senior consultant taking on this brief: [problem].
Before you recommend anything, ask me the five questions you would need answered to solve this properly - the questions a good consultant asks in the first meeting.
Ask them one at a time, wait for my answer to each, and only then give your recommendation.

WHY IT WORKS The quality ceiling of any answer is the quality of the brief. This flips the dynamic: instead of you guessing what the AI needs, it tells you what is missing - and the questions themselves often reframe the problem.

VARIANT - IN A HURRY “List the additional context you would need for a fully accurate answer, then answer anyway, clearly labelling every assumption you have had to make.”

TRY IT ON “Should we build or buy our next scheduling system?” “How do we reduce staff turnover at the Cork site?”

10. The Five Whys

USE IT WHEN A problem keeps coming back no matter how many times it is “fixed”.

Analyse this recurring problem using the Five Whys: [problem].
Ask why, answer it, and ask why again until you reach a systemic root cause rather than a symptom - typically about five levels down.
Distinguish clearly between the immediate cause, the contributing factors, and the root cause.
Then recommend action at each level: a fix for this week, a control for this quarter, and a system change for this year.

WHY IT WORKS A Lean classic from the Toyota tradition, upgraded: the three-horizon ending stops root-cause analysis from producing either a sticking plaster or an impossibly large transformation - and nothing in between.

TRY IT ON *The line stoppage that has happened four times this quarter; invoices that keep going out wrong; the audit finding that reappears every year.*

11. Three Lenses

USE IT WHEN A problem is being framed as purely technical, purely a people issue, or purely a process issue - and you suspect it is not.

Analyse [problem] through three lenses:
Operations - process, flow, waste, capacity.
People - skills, culture, communication, incentives.
Technology - tools, data, automation, integration.
Give findings and one concrete recommendation per lens. Then tell me which lens holds the binding constraint - and why fixing the other two alone would fail.

WHY IT WORKS Most stalled initiatives are stuck in a lens nobody is looking through.

PRO TIP Add a fourth lens for your world: Quality & Regulatory in pharma and med-tech, Food Safety in food, HSQE in engineering, Client Experience in services.

TRY IT ON “The CMMS is fully implemented and maintenance is still reactive.” “Everyone was trained on the new system and adoption sits at 30%.”

DECIDE Commit with clear eyes

A good decision knows its trade-offs, its failure modes, and what would change its mind. These three prompts supply all three - before you commit, not after.

12. The Decision Partner

USE IT WHEN You are choosing between real options and want more than a pros-and-cons list.

Act as a senior decision advisor. We are choosing between: [options].
First, propose the five criteria that should matter most for a business like ours, and weight them.

Then score each option against each criterion, showing your working and the trade-offs honestly, and recommend the strongest choice.

Finish with two things: the conditions under which your recommendation would flip, and the cheapest fast test we could run to find out.

WHY IT WORKS This moves AI from information provider to decision support. The flip-conditions keep the recommendation honest; the cheap test makes it actionable this week.

TRY IT ON *Choosing between two co-packers; lease versus buy on the new CNC machine; which of three sites gets the automation pilot.*

13. The Pre-Mortem

USE IT WHEN Before any significant commitment - a project, a launch, a rollout, a key hire.

Assume it is twelve months from now and [project or plan] has failed - badly enough that we wish we had never started.

Conduct a pre-mortem. Identify the five most likely reasons it failed, the early-warning sign we would have missed for each, and the specific action we could take now to prevent it.

Rank the risks by how much they should worry us today, and tell me which single prevention buys the most protection for the least effort.

WHY IT WORKS Prospective hindsight - imagining the failure as already real - is one of the best-evidenced moves in decision science (Gary Klein's pre-mortem). It surfaces risks that "what could go wrong?" never does, and the early-warning signs turn a worry list into a monitoring plan.

VARIANT - BLACK SWAN STRESS TEST For plans exposed to the outside world, add: "Now generate three highly unlikely but severe external shocks that could hit our sector. For each, show exactly how it breaks this plan, and one adaptation that would make the plan resilient if built in now." (After Nassim Taleb's term for rare, high-impact events.)

TRY IT ON *The ERP go-live. The new production line. The key account renewal you are "certain" about.*

14. Backcasting

USE IT WHEN Planning something ambitious, where forecasting forward from today produces only cautious increments.

Imagine it is three years from now and [project or goal] has succeeded beyond what we thought possible.

Work backwards from that success as if writing its history. What was the critical decision in year two? What bottleneck did we clear in year one that others would not have? What did we deliberately say no to?

Build the reverse timeline of moves that made it happen - then convert it into the three moves we should make this quarter.

WHY IT WORKS Working backwards from success licenses the AI - and you - to propose the non-obvious moves a forward plan filters out, while the final step drags the vision back to this quarter's reality.

TRY IT ON *"By 2029 we are the reference supplier for sustainable packaging in Ireland." "In three years, 40% of revenue comes from services rather than product."*

TRUST Know when to rely on the answer

Confidence is a style, not a credential - an AI sounds equally certain/plausible when it knows and when it guesses. These four prompts make the difference visible, so you know exactly what to verify before you act - and the last two turn the AI loose on its own work to find the gaps before a colleague or client does. They are the heart of practical AI literacy.

15. Evidence vs Opinion

USE IT WHEN An AI answer sounds authoritative and you are about to act on it.

Review your answer above and classify every significant claim as one of: established fact, reasonable inference, assumption, or speculation. Present the result as a table. Then list the claims we would need to verify independently before acting and suggest how to verify each one - a source to check, a person to ask, or a number to pull from our own data.

WHY IT WORKS It makes the difference between knowledge and fluent guesswork visible on one page - and hands you a verification to-do list instead of a vague unease.

TRY IT ON *Any market-size estimate, competitor claim, regulatory interpretation or technical specification, metrics, dates, citations an AI has just given you.*

16. The Confidence Audit

USE IT WHEN A recommendation or plan is on the table and you need to know where the AI is guessing about your business.

Perform a confidence audit on your recommendations.
List your top three conclusions and give each a confidence score from 0-100%.
For each, state exactly what you are assuming about our business, and what specific missing information would move the score close to certainty.
Flag the single assumption that, if wrong, does the most damage.

WHY IT WORKS Watching an AI mark down its own confidence and be realistic.

TRY IT ON *As the final step after any Decision Partner, Pre-Mortem or Conductor output.*

17. The Red Team

USE IT WHEN An AI has produced something you are about to rely on - a plan, a proposal, a piece of analysis - and you want to find the holes before someone else does.

Act as a red team whose only job is to break the output above. Assume it will be read by a sceptical board member, a competitor and an auditor.
Find the gaps: what is missing, what is asserted without support, and what has been quietly assumed.
Find the vulnerabilities: where would it fail under pressure, be misread, or be challenged - and what is the single weakest point?
Rank every issue by severity (critical / significant / minor), state what a critic would say about each, and propose a fix. Then rewrite the three weakest sections.

WHY IT WORKS An AI will argue against its own work as readily as for it - and because it has no reputation to defend, it does so without flinching. Red-teaming turns a polished answer into an inventory of what would embarrass you in a meeting, while there is still time to fix it.

TRY IT ON *A tender response before submission; a business case before the board; any customer-facing document an AI has drafted.*

18. The Output SWOT

USE IT WHEN You want a fast, structured read on whether an AI's recommendation stands up - not just what it gets right.

Run a SWOT on the recommendation above, treating it as a proposal we are about to adopt.
Strengths: what is well-supported, and where it genuinely fits our business.
Weaknesses: what is thin, generic, unsupported or missing.
Opportunities: what it under-exploits, or what becomes possible if it works.
Threats: what could derail it - inside the business or outside it - and what it fails to anticipate.
Present it as a 2x2 grid with three to five points per quadrant, then give one sentence on whether we should adopt, adapt or abandon it.

WHY IT WORKS SWOT is a framework everyone in the room already knows, so the AI's self-assessment arrives in a shape a management team can discuss immediately. Pointing it at the output rather than the business turns a familiar tool into a quality check - and "adopt, adapt or abandon" forces a verdict.

TRY IT ON *The output of any Decision Partner, Backcasting or Conductor prompt; a strategy an AI has just drafted; a supplier proposal you have pasted in for review.*

ACT Turn insight into Monday-morning motion

Analysis is not the deliverable. Most organisations do not have a data problem (maybe access to complete high-quality data maybe); they have an insight-to-action problem. These three prompts are the bridge.

19. So What?

USE IT WHEN You have findings, dashboards or analysis - and a meeting coming where someone will ask "so what do we do?"

Here is a set of findings: [paste findings or analysis].
For every important finding, answer three questions in sequence:
So what does this actually mean?
Why does it matter to this business - in money, time, risk or customers?
What should we do differently as a result?
Rank the implications by importance, and end with the top three actions - not the top three observations.

WHY IT WORKS It refuses to let analysis end at description. The last line matters: actions, not observations, is the discipline most reporting never applies.

TRY IT ON *The monthly KPI pack. The customer-survey results. The energy-usage report gathering dust.*

20. The Executive Brief

USE IT WHEN Good analysis needs to move a decision-maker - in ninety seconds.

Turn everything above into a one-page executive brief:

The key issue in two sentences. The three insights that matter. What they mean for this business. Your single recommendation. The three next actions, each with a suggested owner.

Write it so a busy CEO gets the point in ninety seconds - plain language, no hedging, numbers wherever they exist.

WHY IT WORKS One page and one recommendation force the AI - and you - to decide what matters. It is also the fastest way to make your own analysis presentable upward.

VARIANT - BLUF Write the brief using **BLUF** technique - Bottom Line Up Front, a communication method where you place the main conclusion, decision, or critical takeaway at the very beginning of a message. The busy CEO may get pulled from your presentation 5 minutes in!

TRY IT ON *The output of any other prompt in this playbook.*

21. The Monday Morning Plan

USE IT WHEN A recommendation has been agreed - and the risk is that everyone nods and nothing changes.

Take the recommendation above and turn it into an implementation plan I can start this week. Give me:

1. Three quick wins/ low hanging fruit completable within five working days, each with an owner role and its first concrete step.
2. A 30/60/90-day sequence for the bigger moves, with dependencies flagged.
3. The single most likely point of stall - and how we will know by Friday whether we are on track.

Assume no new budget and no new hires unless you flag one as genuinely essential.

WHY IT WORKS Plans fail in the gap between “agreed” and “started”. This closes it: by the time you finish reading, someone has a first step with their name on it and a Friday check already scheduled.

TRY IT ON *Immediately after any Decision Partner or Executive Brief output.*

ADVANCED A first taste of agentic AI

Everything so far has been one AI, one role. The next step is to make it several: instead of asking for an answer, you instruct the AI to stand up a small internal team of specialists, have them work the problem from different angles, and hand you only the coordinated result. This is, in miniature, exactly how agentic AI systems operate - a controller decomposing a job across specialist agents.

22. The Conductor

USE IT WHEN The problem genuinely needs a team - several disciplines at once - and you have one AI.

TASK: Solve a critical supply-chain disruption.

CONTEXT: Primary European supplier delayed by 3 weeks. Line 2 shuts down in 4 days. The penalty is €50,000 per day.

INSTRUCTION: Stand up three internal specialist agents:

1. <Logistics_Agent> - propose three aggressive, non-standard rerouting options (expedited hot-shot freight, split shipments, alternative modes).
2. <Lean_Operations_Agent> - calculate how to adjust Line 2's cycle time and shift pattern to stretch four days of safety stock to seven, without a full stop.
3. <Risk_Agent> - as a sceptical commercial lawyer, assess the supplier contract for remedies, and draft a firm but collaborative email asking the supplier to cover the air-freight difference.

Have them work the problem, resolve conflicts between their proposals, and present only the final coordinated action plan.

WHY IT WORKS Complex problems are decomposition problems. This shows the AI can be the decomposer, not just the answerer - one prompt returning logistics tactics, scheduling arithmetic and a ready-to-send email, already reconciled.

BEFORE ANYONE ACTS However good the plan looks, finish with 16. and 17. The Confidence Audit and Red team.

Prompt chains - for bigger jobs

The prompts compound. Paste the output of one into the next, and a sequence becomes a workflow. Three proven approaches:

THE DECISION SPRINT Ask Me First → Decision Partner → Pre-Mortem → Executive Brief.

From fuzzy question to defensible decision in one sitting. You finish with a scored choice, a risk plan with early-warning signs, and a one-pager for whoever needs convincing.

THE IMPROVEMENT LOOP Five Whys → Three Lenses → So What? → Monday Morning Plan.

For recurring operational pain: root cause, a full-system view, the business implications, and a first step with someone's name on it by Friday.

THE DEEP DIVE Fast Track → TED Team → Evidence vs Opinion → Executive Brief.

Research you can stand over: you know the field, you know where experts disagree, you know exactly what has been verified - and it fits on one page.

Using GEN AI with confidence

The prompts in this playbook make GEN AI dramatically more useful. Six working practices keep it trustworthy:

1. Treat every output as a first draft from a bright junior colleague – fast and unverified until you check it.
2. Verify before you act. Facts, figures, legal interpretations and anything customer-facing get checked against a primary source. Prompts 15, 16, 17, 18 exist to make that fast.
3. Mind your data. Know which assistant your organisation has approved and what it permits - an enterprise deployment covered by your company's agreements is a different proposition from a public chatbot, and the same prompt may be acceptable in one and out of bounds in the other. Personal data, client-confidential material, commercially sensitive figures and anything under NDA stay out unless your policy explicitly allows them. When in doubt, generalise rather than abandoning the prompt: you will lose far less answer quality than you expect.
4. Keep a named human accountable for every decision AI informs. Under the EU AI Act, organisations already carry AI-literacy obligations for staff using these tools.
5. Date-stamp your research. Models have knowledge cut-off dates. For anything time-sensitive - prices, regulations, standards, market data - ask the tool what may be out of date and confirm current facts from a live source.
6. Build a prompt library. When a prompt works on your problems, save it, share it, improve it. A shared library becomes a genuine company asset - this playbook is the starter kit.

Your first week

Five days, one prompt a day, on real work - the fastest way to feel the difference between using AI and working with it.

Day	Prompt	Do this
Mon	9. Ask Me First	Bring it a real problem you are facing this week - and answer its five questions properly.
Tue	10. The Five Whys	Point it at the niggle that keeps coming back.
Wed	13. The Pre-Mortem	Run it on the project you are most confident about.
Thu	15. Evidence vs Opinion	Run it on something an AI told you recently, and see what survives.
Fri	20. The Executive Brief	Turn the week's best output into one page and share it with a colleague.

And whenever an answer disappoints, run the anatomy check - Role, Context, Objective, Constraints, Criteria, Format, Challenge. The missing element is usually the fix.

The next page is designed to be printed and pinned up beside your desk.

Quick reference - PIN UP

#	Prompt	Reach for it when...
1	The Mom Test	You have a product idea and need to know the problem is real before building.
2	The Fast Track	You need to get up to speed with specialists, fast.
3	The TED Team	You want depth and the case against - not a definition.
4	Make It Transferable	An idea from elsewhere needs translating into your sector.
5	The 10th Man	Everyone agrees - and that worries you. (It should.)
6	The Fair Trial	A big idea needs judging, not flattering.
7	Fresh Eyes	A familiar process needs an outsider's obsessions.
8	First Principles	"How it's done" might just be habit.
9	Ask Me First	The problem is important and fuzzy.
10	The Five Whys	The same problem keeps coming back.
11	Three Lenses	The framing feels too narrow - ops vs people vs tech.
12	The Decision Partner	You are choosing between real options.
13	The Pre-Mortem	You are about to commit to something significant.
14	Backcasting	Ambition needs a plan, not a forecast.
15	Evidence vs Opinion	An answer sounds confident and you are about to act on it.
16	The Confidence Audit	You need to know where the AI is guessing about your business.
17	The Red Team	You need the holes found before someone else finds them.
18	The Output SWOT	A recommendation needs a verdict: adopt, adapt or abandon.
19	So What?	Findings need to become implications and actions.
20	The Executive Brief	A decision-maker has ninety seconds.
21	The Monday Morning Plan	Agreement risks becoming inertia.
22	The Conductor	The problem needs a team - and you have one AI assistant.

About Clarke Analytics

Clarke Analytics is an Irish AI and data analytics consultancy based in Mallow, Co. Cork, working with organisations from small medium enterprises to multinationals on practical data analytics, AI adoption and governance-first AI.

Dave Clarke has led data analytics training for ICBE members since 2021. This playbook accompanies the ICBE webinar 'Why Before How Getting Value from AI' and the ICBE Advanced Productivity Skillnet programme

From AI Curiosity to AI Capability 2026

Leading Digital Transformation: A practical cross-functional programme in Data Analytics, AI and Agentic AI

KEEP THE CONVERSATION GOING WITH

dave@clarkeanalytics.com · www.clarkeanalytics.com

About ICBE Advanced Productivity Skillnet

This comprehensive Skillnet is committed to promoting innovation and productivity through industry-led training for the manufacturing and service sectors in Ireland. Advanced methodologies and technology are key focus areas.

Focusing initially on Lean, Six Sigma and Shingo, this Skillnet has expanded to include key technologies and methodologies from automation to data analytics and additive manufacturing.

KEEP THE CONVERSATION GOING WITH

aidan@icbe.ie · www.icbe.ie

About ICBE Business Excellence Skillnet

The objective of ICBE Business Excellence Skillnet is to develop the knowledge, skills and qualifications of employees by increasing access to skills development opportunities and encouraging ongoing professional development.

The strategic purpose is to focus on providing efficient and effective enterprise-led training and networking opportunities for companies to develop their workforce skills and talent so they are prepared for the political, technological and competitive risks that are already here and for the future.

KEEP THE CONVERSATION GOING WITH

grainne@icbe.ie · www.icbe.ie

© 2026 Clarke Analytics Ltd. Prepared for the member companies of the Irish Centre for Business Excellence Skillnets, who are welcome to use and share this guide within their organisations.