



Case study

AI-assisted macro planning

How a chip-development team compressed multi-year program planning from weeks to hours, and how the team's iteration carried the output the rest of the way.

FTTM methodology series.

A case study from a semiconductor program. The team used a small set of AI tools and a deliberate iteration discipline to produce executive-grade planning artifacts in a fraction of the usual time. AI does about 80%. The team scrubs the rest.

Prepared by
lateralworks
Case study

Date
May 2026
Rev 1.1

Online
lateralworks.com
FTTM series

Table of contents

AI-assisted macro planning

Abstract	03
01 The setup	04
02 The three-step flow	06
03 The iteration loop	10
Working principle: roughly right, fast	14
04 What the research says	15
05 What we learned	18
References	22

Core thesis. The model that works compresses about 80% of the synthesis work into hours and routes the remaining 20% (judgment, calibration, local context) through the team. The team is not validating the AI. The team is building the plan, and AI is helping it get there faster.

Overview

Abstract

Building a multi-year program plan used to take weeks of meetings and drafts and reviews. A cross-functional team would gather inputs from documents, transcripts, and emails, argue through the structure, iterate the artifacts by hand, and converge on a charter, a milestone matrix, and a detailed schedule. The cycle was painful and often produced something too rigid to revise. By the time the plan was published, parts of it were already stale.

This paper is a case study from a semiconductor program that replaced that cycle with an AI-assisted three-step flow. A senior project manager fed source documents and meeting transcripts into a large language model and asked it to produce a milestone timeline, then a milestone matrix, then a macro schedule that could be imported into a project-management tool. Each step took hours. The team scrubbed and refined each draft, fed the corrections back into the next pass, and used the time saved to argue about the program instead of about the artifact. The first complete plan was on the table in three days. The conventional cycle would have taken three to four weeks.

The case study documents the flow, the prompts that worked, the mistakes that surfaced, and the human disciplines that made the output trustworthy. It places the working pattern inside the broader research literature on AI-augmented knowledge work: the Harvard / BCG study of 758 consultants on the "jagged technological frontier" [1], the Stanford / MIT evidence on customer-support productivity [2], and the MIT writing experiment in *Science* [3]. The three studies converge on a single finding. AI shortens the time to a usable draft sharply, and the value of the final output depends on what the team does with the draft.

The paper closes with operating principles any program team can adopt. The flow itself does not depend on the chip context, and teams in clinical research, software, and integration programs have reported similar results.

Section two presents the flow itself with the three artifacts it produces. Section three describes the iteration loop that ties the three steps together. Section four places the case-study findings inside the broader research literature, and section five distills nine operating principles. The paper is short. The case is simple, and the discipline behind it is simple too.

01

Section one The setup

A semiconductor program had reached the point where the team could no longer plan the way it had before. The product was a multi-year, multi-partner stacked-memory device. The schedule depended on five process and design milestones, each gated by partner decisions and each carrying its own doneness criteria, assumptions, and risks. The core team was eight people; peak engineering headcount would run into the hundreds; the budget would land in the tens of millions. The conventional planning cycle would have taken weeks of meetings to produce a usable draft, and the draft would have been obsolete the moment a partner decision shifted.

Three strong inputs were already in hand. The product owner had produced an early milestone timeline at quarter-level granularity. The technical program manager had run a structured milestone scrub with the fab/process, design, and integration leads, and the meeting had been transcribed. A risk register was already maintained in a spreadsheet. Scattered emails and chat threads also held individual decisions that had never been merged into the master plan.

The traditional path forward would have assigned a planner to merge those inputs by hand into a project charter, then circulated the charter for comment, then incorporated the comments, then converted the result into a Microsoft Project schedule. The team had run that cycle before and abandoned it. It was too slow to keep up with the program and too brittle to absorb partner decisions, and it always pinned the work on a single planner who became the bottleneck.

What the team had

Inputs and constraints

A working list of what fed the planning process at the start.

Inputs

Milestone scrub minutes (PDF). An eight-page meeting record capturing technical discussion across process, design, and integration. Decisions, owners, open unknowns, and a top-10 milestone set rebuilt from scratch in the room.

Risk register (spreadsheet). Twenty-five tracked risks at three severity levels with named owners and first-mitigation steps.

Otter.ai meeting transcript (text). A verbatim record of the milestone scrub, used by the AI to recover the "why" behind decisions that the formal minutes had compressed away.

Email threads and chat comments. Subsequent corrections, partner-trip readouts, and individual decisions recorded outside the formal documents. Each of these would, in conventional planning, have to be hunted down and merged manually.

Project charter draft (PDF). An eight-page document carrying mission, success criteria, assumptions, risks, and an early milestone timeline at quarter-level granularity. Strong on framing, weak on sub-task structure. Generated last in the input chain because the charter inherits structure from the prior artifacts.

Constraints

Time pressure. The next executive review was three weeks away, and a partner-decision trip was scheduled for two weeks after that. The plan had to be coherent enough to support both events.

Distributed team. Nine named contributors across three continents, none of whom could afford a multi-day planning offsite.

Partner dependencies still open. Three of the five major milestones depended on partner-side decisions that would not close until the trip. The plan had to be structured to accept those decisions cleanly when they arrived rather than to assume them.

Decision

The program manager proposed an AI-assisted approach: feed the documents and transcripts into a large language model, ask it to produce three connected artifacts (a milestone timeline, a milestone matrix, and a macro schedule), and use the team's time to scrub each draft rather than to build it. The first iteration would be rough. The second would be much closer. The third would be publishable. The total elapsed time would be days, not weeks. The bet was simple: AI gets the team to roughly right very fast, and the team's time — the scarce resource — is best spent on judgment, not on synthesis.

02 Section two

The three-step flow

The flow is three connected steps, and each step produces a different artifact for a different audience. The milestone timeline is for the executive sponsor: a one-page view that fits on a charter page and answers the question "what is this program, what does it look like, what could kill it?" The milestone matrix is for the team: a working table with doneness criteria, sub-tasks, assumptions, and linked risks for every milestone. The macro schedule is for the project-management tool that will eventually drive weekly execution.

Each artifact builds on the one before it. The matrix inherits structure from the timeline and adds operational detail. The schedule inherits structure from the matrix and adds calendar dates and dependencies. Corrections made at any level flow into the next pass without rebuilding from scratch.

AI generates the draft. The team scrubs and refines. The schedule is the product.

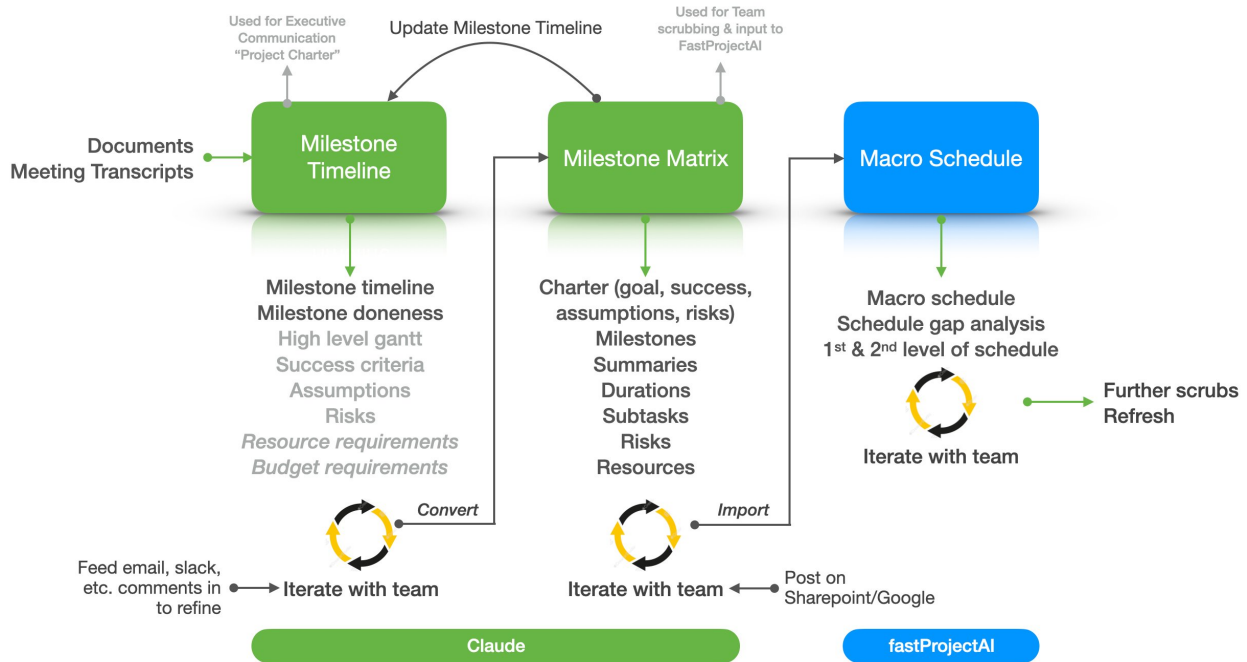


Figure 1. The three-step AI-assisted macro planning flow. Each step produces a draft, the team scrubs each draft, and comments feed back into the next iteration.

Step one

Milestone timeline

The first step gathers and synthesizes. The AI receives every available source: documents, transcripts, the risk register, side conversations. It is asked to produce a one-page milestone timeline in a fixed format. A horizontal time axis, one to nine milestones plotted against it, three columns underneath for success criteria, assumptions, and risks. The format is the contract. The AI fills it in.

The first draft is rough. Some milestones are mis-named. Some doneness criteria belong on a different milestone. The risk column over-weights items that were prominent in the most recent meeting and under-weights items that have been quietly important for longer. These are the visible errors. The team scrubs them in a 30-minute meeting and the program manager re-prompts the AI with the corrections. The second iteration is materially better. The third is publishable.

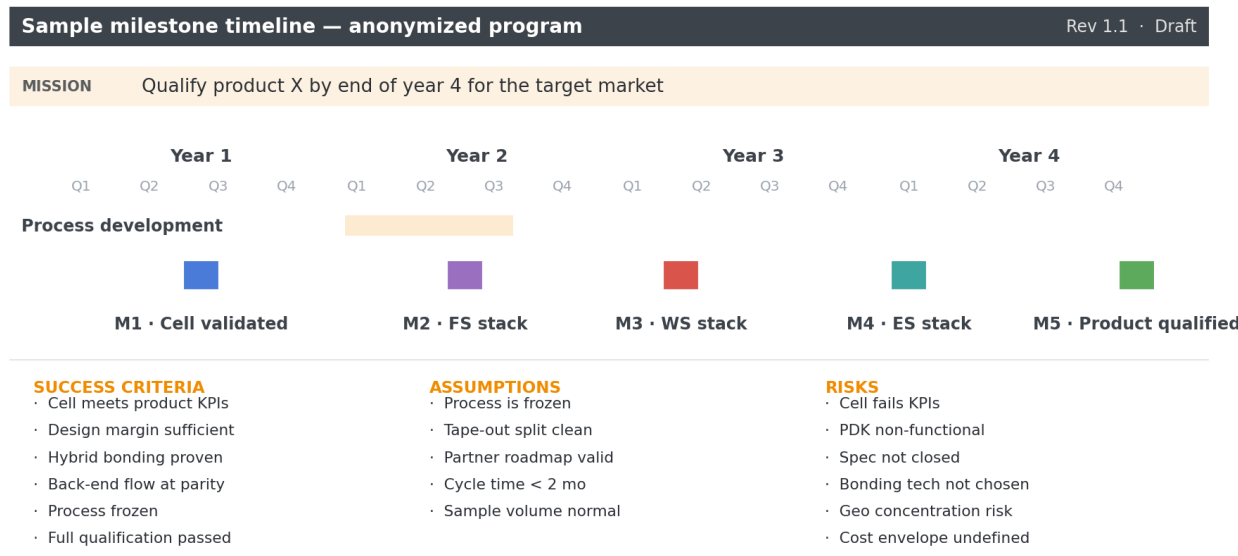


Figure 2. A representative milestone timeline output, anonymized. The AI produces this structure on the first prompt; the team's iteration adjusts wording, severity, and milestone scope rather than the layout itself.

The output becomes the executive page of the project charter. It is the artifact the sponsor reads and the artifact the team uses to test whether the program structure is coherent. If the timeline cannot be drawn cleanly on one page, the milestones are wrong.

Observation. The first AI draft of the timeline produced eleven milestones. The team scrubbed it down to five. The reduction did not require new information. It required a human judgment about which gates were load-bearing and which were intermediate readouts that belonged inside other milestones.

Step two

Milestone matrix

The second step distills and normalizes. The AI takes the agreed milestone timeline and expands each milestone into a row of a working table. The columns are fixed: milestone identifier, summary task, milestone name, doneness criteria, key sub-tasks, assumptions, and risks. Same format on every row. Same structure means downstream tools can ingest the table directly.

The matrix is where the program shifts from executive language to execution language. Doneness criteria become unambiguous. "1st Si 3D stack delivered: hybrid bonding selected, stacking partner chosen, stack sorted and tested" replaces "stack working." Each sub-task carries an implicit owner. Assumptions become testable. Risks become mappable to the existing risk register.

Sample milestone matrix — anonymized program

M#	Summary task	Milestone name	Doneness criteria	Key subtasks	Assumptions	Risks
M1	Cell development	Cell validated	SL vehicle done · cell meets KPIs · integration path set	Tapeout SL · monthly KPI review · DOE for cap height	jDP in place · SL operational by Q2	Cell KPI fail · spec slipped
M2	FS stack	FS 3D stack validated	Base + memory die delivered · 1st Si stack · stack validated	Partner trip · base tapeout · stacking · sort & test	Bonding decision closed · partners signed	PDK broken · bonding not picked · async dies
M3	WS stack	WS stack validated	Retrofit dies · 2nd Si validated · variation budget set	FS learning distill · retrofit freeze · 2nd Si tapeout	Learnings incremental · cycle time < 2 mo	Architectural rework · stack-level variation
M4	ES stack	ES stack validated	Eng samples on prod-intent · JEDEC compliance · system val	Sample qty · prod-intent flow frozen	No respin needed · datasheet draft · ship approval	Sample latency · respin between WS and ES
M5	Qualify	Product qualified, released	Qual passed · retention met · endurance · UBER	Qual program per RCM · char test · stocking plan	No respin · ELFR after release	Respin needed · margin insufficient

Figure 3. A representative milestone matrix, anonymized. The matrix is the bridge between executive framing and team execution, and it is the artifact most heavily scrubbed since every cell is read by the responsible engineer.

The matrix is scrubbed twice. The first scrub is by the program manager and the AI. The AI is asked to flag rows where the doneness criteria do not match the milestone name, where assumptions are duplicated across rows, where risks are mapped to the wrong milestone. The second scrub is by the team. Each engineer reads the rows in their domain and corrects them in place. Both scrubs are cheap. Together they take the matrix from 80% to 95%.

Observation. The fastest part of the entire planning cycle was the matrix. Once the timeline was settled, the AI generated the matrix in a single prompt, and the team's scrub was crisp because the structure was already familiar. Time to publishable: under one working day.

Step three

Macro schedule

The third step organizes and schedules. The AI converts the matrix into a structured input that a project-management tool can ingest. In the case-study program, the target was a desktop scheduling tool with built-in milestone visualization. Microsoft Project, Smartsheet, or Asana would work the same way. The artifact is no longer free-text. It is a machine-readable plan.

The schedule adds calendar dates and dependency arrows. The AI estimates durations from the sub-tasks and the matrix's effort hints. The team scrubs the durations. This is the only step where the AI consistently mis-estimates, because it has no calibration to the team's actual fab cycle time, partner cadence, or holiday calendar. The team adjusts. The dates snap into place.

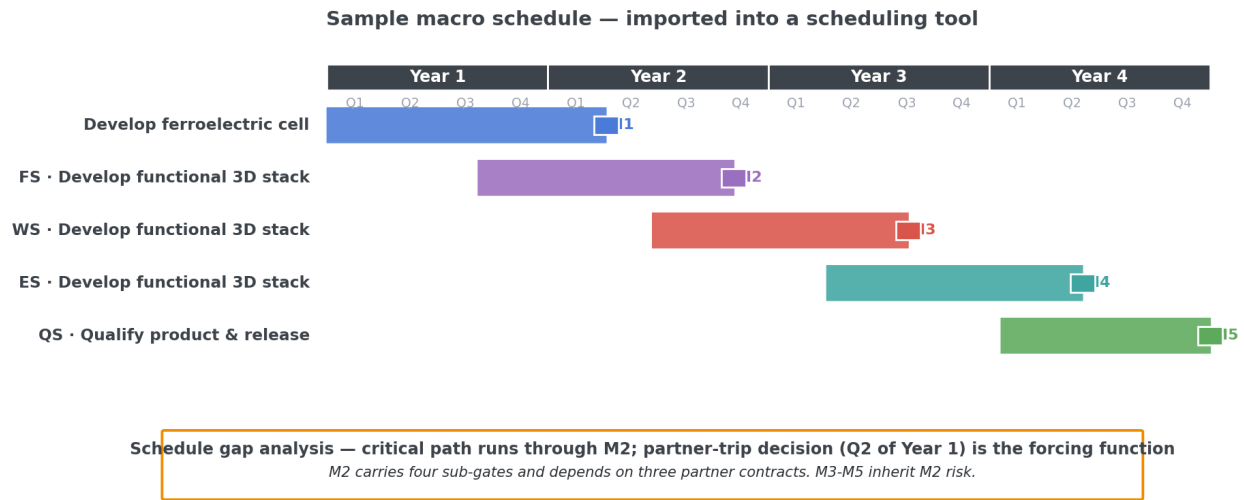


Figure 4. A representative macro schedule view, anonymized. The critical-path callout below the gantt highlights the gating decision the executive review needs to see.

The schedule surfaces what the prior two artifacts cannot: the critical path. Once milestones, sub-tasks, and dependencies are in the tool, the gating sequence becomes visible. In the case-study program, the second milestone (M2) carried four sub-gates and three partner contracts. M3, M4, and M5 all inherited risk from it. The schedule made that visible in a way the timeline and matrix could not.

What the team did with the time saved. Instead of arguing about the format of the plan, the team argued about the program. Whether the partner trip should drive the M2 close or follow it. Whether the base die should be a separate track. Whether the qualification gate should be re-weighted. Those arguments produce a better program. They were possible because the artifacts arrived in days instead of weeks.

03

Section three

The iteration loop

The three-step flow describes the artifacts. The iteration loop describes how each artifact is brought from rough first draft to publishable. The loop has six stations: gather, distill, synthesize, organize, scrub, refine. The first three are AI-driven. The last three are team-driven. The loop closes when the team's corrections become the prompt for the next AI pass.

No iteration happens in isolation. Every iteration is informed by the previous one and feeds the next. The discipline that distinguishes a useful AI-assisted plan from a mediocre one is to not stop after the first draft.

The loop

Six stations, two roles

The six stations are not equal. Three are nearly free; the AI does the work in seconds. Three require human attention and judgment and cannot be shortened. Recognizing which is which is the first move.

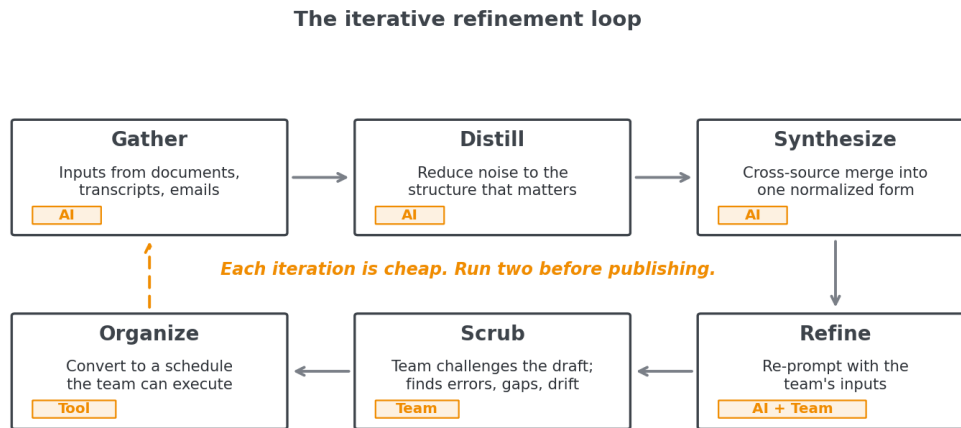


Figure 5. The six-station iteration loop. The top row is AI-cheap, the bottom row is human-essential. The dashed loopback arrow is the discipline that makes the system work.

Gather, distill, synthesize — AI

These three stations are where AI dominates. **Gather** means reading the inputs. A 30-page PDF, a 90-minute transcript, three email threads, a risk spreadsheet, and a slide deck do not need to be read by the planner anymore. The AI ingests all of them in seconds. **Distill** reduces the noise. Long documents carry far more text than is needed to plan against, and the AI is unusually good at compressing the structural content while preserving what matters. **Synthesize** is the cross-source merge. When the same milestone is described differently in the charter, the transcript, and an email thread, the AI normalizes them to a single representation. This is the work that used to take a planner days. It now takes minutes.

Organize, scrub, refine — humans

Organize is partly tool-driven. Importing the matrix into a scheduling tool, mapping fields, validating dates. **Scrub** is the work that has to be human. The team reads the AI's output looking for three classes of error. Factual mistakes (wrong dates, wrong owners, wrong numbers). Structural drift (a milestone that has subtly changed scope across iterations). Missing local context (an assumption the team knows is no longer valid but the AI cannot have known). **Refine** closes the loop. The team's corrections become the input prompt for the next AI pass.

Iteration discipline

Three is enough

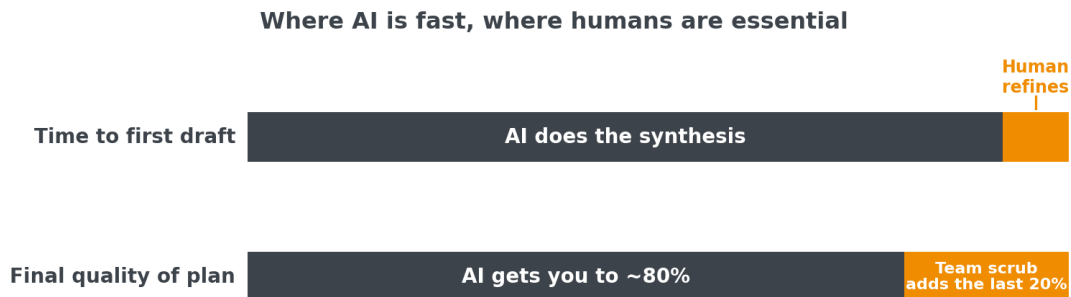
**"One iteration
is not enough.
Three is usually enough."**

First draft rough. Second meaningfully better. Third publishable.
After the third, returns shrink fast.

The split

Where AI is fast, where humans are essential

The most-quoted line from this case study is the simplest. AI gets you to 80%. Humans add the last 20%. The shape of the split is what makes the pattern work. AI is fast at the first 80% (the synthesis, the structuring, the rough draft) and slow at the last 20% (the judgment and calibration that depend on context the AI does not have).



The order matters. AI first cuts the time-to-draft from weeks to hours. Then humans add the judgment that AI cannot supply.

Figure 6. The 80 / 20 split. AI compresses time-to-draft sharply, and the team scrub determines final quality.

The order matters. AI first, humans second. Reverse the order, with humans drafting and AI polishing, and the gain shrinks because the bottleneck (synthesis) was never compressed. It also breaks the iteration loop. A hand-drafted artifact is much harder to re-prompt with a structured correction than an AI-drafted one, because the structure itself is attached to the author.

The Pareto principle is the right reference here, in its operational form: roughly 80% of the value comes from roughly 20% of the work [4]. In AI-assisted planning, the inversion is what creates the time savings. The 80% of the work that produces the basic structure becomes nearly free, and the team can spend its full attention on the 20% that determines whether the plan is correct.

Observation. A frequent failure mode is for a team to demand 100% AI accuracy before publishing the first draft. That is the wrong target. The first draft only needs to be good enough that the team can scrub it efficiently. Polishing the AI to 100% is what the second and third iterations are for.

Working principle

Roughly right, fast

**"It is better to be
roughly right
than exactly wrong."**

A working principle for AI-assisted planning
lateralworks Fast Time To Market (FTTM) methodology

04

Section four

What the research says

The case study is one data point. The broader research literature on AI-augmented knowledge work has produced converging evidence across very different settings (management consulting, customer support, professional writing). The same pattern shows up in each: substantial productivity gains on tasks that fall inside the AI's capability frontier, paired with serious risks on tasks that fall outside it. The case-study flow is one practical answer to the question of how to capture the gains while avoiding the risks.

The three studies summarized below are the strongest peer-reviewed evidence available as of early 2026. None is specific to program planning. Each illuminates something the case study assumed without proving: that AI compresses time to draft, that the gains are uneven across tasks, and that human iteration determines final quality.

The studies

What three field experiments found

Harvard / BCG — 758 consultants

The 2023 Dell'Acqua et al. field experiment with Boston Consulting Group [1] is the most-cited study in this literature. Seven hundred and fifty-eight BCG consultants, about 7% of the firm's individual-contributor workforce, were randomly assigned to one of three conditions: no AI, GPT-4 access, or GPT-4 access plus prompt-engineering training. They were given eighteen realistic consulting tasks and graded on speed and quality.

Inside the AI's capability frontier, consultants with AI completed 12.2% more tasks, took 25.1% less time, and produced output 40% higher in quality than the control group. Outside the frontier, on tasks that look superficially similar but fall just beyond what the AI can reliably do, consultants with AI were 19 percentage points more likely to produce wrong answers than consultants without AI [1]. The study coined the phrase "jagged technological frontier" for this property.

Stanford / MIT — 5,179 support agents

The Brynjolfsson, Li, and Raymond study, first released as NBER Working Paper 31161 in 2023 and published in the *Quarterly Journal of Economics* in 2025 [2], examined the staggered introduction of a generative-AI conversational assistant across 5,179 customer-support agents at a Fortune 500 company. Access to the AI tool increased productivity, measured as issues resolved per hour, by 14% on average.

The unequal distribution of the gains carries the planning lesson. Novice and low-skilled agents saw a 34% productivity improvement. Experienced and high-skilled agents saw a minimal effect. The authors interpret this as the AI capturing the best-practice patterns of the most-productive agents and disseminating them to everyone else [2]. Translated to planning: the AI does the synthesis the way an experienced planner would, even when the user is not yet experienced.

MIT — 453 professional writers

The Noy and Zhang experiment, published in *Science* in July 2023 [3], assigned occupation-specific writing tasks to 453 college-educated professionals (managers, HR specialists, grant writers, marketers, consultants, data analysts). Half were randomly given access to ChatGPT. Average time on task fell by 40%. Output quality rose by 18%. Variance between high-ability and low-ability writers compressed; the gain to lower-ability writers was substantially larger.

The result that matters most for planning work is the change in how people worked. ChatGPT did not just speed up rough drafting. It shifted the structure of the task. The treatment group spent less time on rough drafting and more time on idea generation and editing [3]. That is the same shift the case study describes. The team stopped writing the plan and started judging it.

Two patterns

Centaur and cyborg

Two distinct patterns of human-AI integration have emerged in the broader research. The terms come from chess. In 1998, the year after losing to IBM's Deep Blue, Garry Kasparov organized the first "Advanced Chess" tournament. Human players paired with chess engines competed as teams. The image of the human-on-horse hybrid stuck as the name for the resulting working pattern [5]. The Dell'Acqua team found that consultants in their study clustered into two related patterns. They labeled them *centaur* and *cyborg* [1].

Two patterns of working with AI on knowledge tasks

Both produce gains on tasks inside AI's capability frontier; the integration style determines how those gains compound.

Centaur	Clear human/AI division of labor
<ul style="list-style-type: none"> · Human chooses the overall approach · AI handles the heavy lifting on specific tasks · Boundaries between the two are explicit · Best when the task partitions cleanly 	
Cyborg	Continuous human-AI back-and-forth
<ul style="list-style-type: none"> · Human and AI iterate inside every task · Steady stream of prompts, edits, re-prompts · No clean boundary between roles · Best when the task is ambiguous and evolving 	

Pattern adapted from Dell'Acqua et al. 2023, "Navigating the Jagged Technological Frontier"

Figure 7. The two integration patterns. Both produce gains inside the AI's capability frontier. The difference is how the human and AI exchange work. Pattern adapted from Dell'Acqua et al. 2023 [1].

Centaur users define a clear division of labor. The human chooses the approach, the AI handles specific tasks within that approach, and the boundary between the two is explicit. **Cyborg** users blur the boundary. They iterate inside every task, with a continuous back-and-forth of prompts, edits, and re-prompts. Neither pattern is universally better. The case-study flow is closer to the centaur pattern at the artifact level (AI produces the timeline, the team scrubs it) and closer to the cyborg pattern inside each iteration (the planner re-prompts continuously while reviewing).

The choice is task-dependent. Centaur is better when the task partitions cleanly: a well-defined sub-task the AI can take entirely, with the rest as human work. Cyborg is better when the task is ambiguous and evolving, when the right next step depends on what the AI just produced. Macro planning is mostly the first kind at the artifact level (the timeline can be generated independently and then scrubbed) and mostly the second kind during iteration (each round of corrections changes what the next prompt should ask for).

The takeaway from the literature. *The productivity gain from AI is real and substantial, and it is not uniform. It depends on whether the task lies inside the AI's capability frontier, on how the human integrates with the AI, and on whether iteration discipline is enforced. The case-study flow is one concrete answer.*

05

Section five

What we learned

Nine working principles came out of the case study. Each is phrased as a discipline rather than a tip, because each one is easy to ignore and the cost of ignoring it is large. Together they describe the difference between a team that captures the AI productivity gain and a team that does not.

The principles are portable. None is specific to semiconductor development. A software product team, a clinical research program, a construction project, a corporate restructuring all face the same planning problem and can use the same flow.

Nine principles

How to operate

1. Always ask for a summary

The first artifact in any AI-assisted planning cycle is a summary or abstract of the source material. It is the cheapest way to verify the AI has read the inputs accurately. If the summary is wrong, the plan will be wrong. Catching it at the summary stage costs minutes. Catching it at the schedule stage costs days.

2. Get to 80% fast, then iterate

The first AI pass should aim for the basic structure and content, not for polish. Use the first pass to find out what the AI got wrong, what is missing, and what is over-weighted. The second pass corrects. The third refines. Trying to land a perfect plan in one prompt is the failure mode that wastes the most time.

3. The more base data, the better the output

AI synthesis quality scales with the volume and structure of inputs. A team that feeds in only the project charter will get a thinner output than a team that feeds in the charter, the meeting transcript, the risk register, and the partner emails. The team's job in input gathering is to be exhaustive. The AI handles the selection.

4. Use every artifact you have

Emails, chat threads, Slack messages, transcripts, slide decks, side spreadsheets. Conventional planning loses these because no one has time to merge them. The AI merges them automatically. Treat every artifact as a candidate input to the next iteration.

5. Tell people the document is a draft

When sharing AI-generated artifacts for review, label them explicitly as drafts and ask for specific corrections. "This is a draft and it needs your inputs. What is wrong? What needs to be changed?" That invites the reviewer to engage as a collaborator instead of as a judge. It also creates investment in the final product, which improves trust in the result.

6. Store the sessions

AI tools that retain context across conversations get smarter about a specific program over time. A planning project should have its own dedicated workspace into which every related session is filed. The fifth iteration of a plan is much better than the first, partly because the AI has accumulated context about the program that no individual session would carry.

7. Always run the humanizer

AI-generated text has stylistic tells (the em dashes, the recurring transitional phrases, the slightly over-formal register) that signal "AI wrote this" to a reader. Before publishing, run the text through a humanizer pass that flattens those tells. The content of the plan is what matters. The voice should be invisible.

8. Identify sources and provide references

When a plan makes a non-obvious claim (a benchmark headcount, a typical schedule duration, a market data point), the AI should be asked to identify the source and surface it in a reference list. The discipline catches hallucinations early. It also makes the plan defensible against challenge.

9. Have a skeptic on the team

A team without a skeptic over-trusts AI output. A team with a skeptic catches the errors the rest of the team would have missed. The skeptic's role is to refuse to accept the AI without scrutiny. The case-study team had a skeptic. The plan is better for it.

Closing

What this means more broadly

The case study has been described as if it were unusual. It is a particular instance of a working pattern that is becoming the default for knowledge work. The pattern compresses synthesis time, surfaces the team's judgment as the scarce resource, and rewards iteration discipline.

Three observations close the paper. First, the pattern is portable across industries. The flow described here was used for a hardware development program. The same flow has produced similar results in customer-success rollouts, clinical study planning, M&A integration playbooks, and regulatory response programs. Nothing about the underlying mechanism depends on the program domain.

Second, the gains are not uniform across team members. The Brynjolfsson, Li, and Raymond evidence [2] is unambiguous on this. AI levels the productivity distribution, with the largest gains going to less-experienced workers. In planning work, that means a junior planner with AI can produce a draft plan that a senior planner without AI would take a week to match. The senior planner's comparative advantage shifts toward judgment and scrutiny, where it should have been all along.

Third, the working principle behind the entire case is older than AI. "It is better to be roughly right than exactly wrong" is the operating discipline of most fast-moving teams. The structural feature of AI-assisted work is that it captures this discipline operationally. AI lets a team be roughly right in hours instead of in weeks. Whether the team uses the time saved to be more right, or simply to ship faster, is the choice that determines whether AI assistance produces lasting value.

The closing principle. AI gives you a fast first draft. The value of the final plan depends on what the team does with the draft. Run two iterations. Scrub the output. Argue about the program, not the artifact.

R

Sources

References

The references below are the primary sources for the empirical claims in this paper. Where a study is available in both a working-paper version and a peer-reviewed version, the citation points to the peer-reviewed version of record but lists the working paper for accessibility.

- [1] Dell'Acqua, F., McFowland III, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K. C., Rajendran, S., Kraymer, L., Candelon, F., and Lakhani, K. R. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality." Harvard Business School Working Paper 24-013, September 2023. Forthcoming in *Organization Science*.
https://www.hbs.edu/ris/Publication%20Files/24-013_d9b45b68-9e74-42d6-a1c6-c72fb70c7282.pdf
- [2] Brynjolfsson, E., Li, D., and Raymond, L. R. "Generative AI at Work." *Quarterly Journal of Economics*, vol. 140, no. 2, 2025, pp. 889–942. NBER Working Paper 31161, April 2023. <https://doi.org/10.3386/w31161>
- [3] Noy, S., and Zhang, W. "Experimental evidence on the productivity effects of generative artificial intelligence." *Science*, vol. 381, no. 6654, July 2023, pp. 187–192. doi:10.1126/science.adh2586
- [4] Pareto, V. *Cours d'Économie Politique*. Lausanne: F. Rouge, 1896–1897. (Source of the original Pareto distribution observation. The operational "80/20 rule" interpretation is attributed to Joseph M. Juran in his 1951 *Quality Control Handbook*.)
- [5] Kasparov, G. *Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins*. PublicAffairs, 2017. (Origin of the "centaur" concept for human-AI collaborative play, first applied in Kasparov's 1998 Advanced Chess tournament.)
- [6] Mollick, E. *Co-Intelligence: Living and Working with AI*. Portfolio, 2024. (Operational guidance on the "jagged frontier" and on iteration discipline as the locus of AI-augmented productivity.)
- [7] McKinsey & Company. "The state of AI in early 2024: Gen AI adoption spikes and starts to generate value." Global Survey on AI, May 2024. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-2024>
- [8] McKinsey & Company. "The state of AI: How organizations are rewiring to capture value." Global Survey on AI, March 2025. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
- [9] Peng, S., Kalliamvakou, E., Cihon, P., and Demirer, M. "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot." arXiv preprint, February 2023. <https://arxiv.org/abs/2302.06590> (Companion evidence on AI compression of time-to-draft for software-development tasks; 55.8% faster task completion in the treated group.)
- [10] Kellogg, K. C., Lifshitz-Assaf, H., Randazzo, S., Mollick, E., Dell'Acqua, F., McFowland III, E., Candelon, F., and Lakhani, K. R. "Don't Expect Juniors to Teach Senior Professionals to Use Generative AI: Emerging Technology Risks and Novice AI Risk Mitigation Tactics." Harvard Business School Working Paper 24-074, March 2024.
- [11] lateralworks. "Fast Time To Market (FTTM) methodology — team best practices." Internal methodology paper, 2026.
- [12] lateralworks. "Internal engagement records." Case-study data drawn from a complex semiconductor program, April–May 2026. Anonymized for this paper. Details on file with the program team.

A

Appendix

Example outputs

The pages that follow are example projects, redacted and genericized, that show what each step in the AI-assisted flow produces in practice.

- A.1 Memory program — Milestone Timeline and Project Charter (Step 1).**
An eight-page project charter built around a Step 1 milestone timeline. Generated by an LLM from program inputs — a brief, the meeting transcript from a milestone scrub, a risk register, and partner emails — and iterated with the core team before sign-off.
- A.2 Project Beta — Milestone Matrix, card form (Step 2).**
A four-page expansion of the Step 2 matrix into project scope plus milestone-by-milestone cards. The card form is what the team uses for review and discussion: each milestone shows doneness criteria, sub-tasks, assumptions, and linked risks side by side.
- A.3 Project Beta — Milestone Matrix, table form (input to Step 3).**
The same matrix collapsed into a single landscape page. This is the form that gets imported into fastProjectAI to generate the macro schedule and conduct critical-path and gap analysis.
- A.4 Fab1 program — Macro schedule with risk and procurement analysis (Steps 1–3 extended).**
A seven-page feasibility-stage planning bundle for a multi-billion-dollar fab program. Shows all three artifacts produced for a single program — milestone timeline (Step 1), milestone-detail cards (Step 2), and macro schedule with critical path (Step 3) — extended with program context, top risks, CAPEX procurement timeline, and a ranked priority mitigation list. Same flow, larger program, longer horizon.

Steps 1 and 2 are generated by LLM models and iterated multiple times with the core team before this data makes its way into the project-management software used to execute the project.

DDR5 NV-DRAM+ — PROJECT CHARTER


Persistent memory for AI scale-out · Foundry Partner 1B 16Gb 3D DDR5 NV-DRAM+ (4H stack)

REV 1.1 · 24 APRIL 2026

Updated schedule — realistic baseline

MISSION

Qualify a Foundry Partner 1B 16Gb 3D DDR5 NV-DRAM+ (4H stack) for AI data-center scale-out by December 2029.



CORE TEAM · Exec sponsor: Exec Sponsor · Product boss (E2E owner): Product Boss · TPM: TPM · Design: Design Lead · Fab/Process: Fab/Process Lead · Device: Device Lead · Post-Fab: Post-Fab Lead · Product: Product Lead

SUCCESS CRITERIA

- Ferroelectric cell meets product KPIs (imprint, polarization, temperature envelope)
- Design margin sufficient to sample at 3200 MT/s across full operating window
- 4H W2W hybrid bonding proven; TSV connectivity verified across stack
- Manufacturing-ready back-end flow proven (bonding + package + test) at DRAM parity on test cost
- Foundry Partner 1B process frozen; base die flow frozen at Logic Foundry + Interface Design Partner
- Full JEDEC DDR5 qualification passed — not by similarity to DDR4
- Retention 3 months @ 85°C; endurance 10¹⁵ cycles; UBER 10²⁴
- Yield-to-target (~80%) plan developed; margin-to-goal plan developed
- Product chip-level functional and performance requirements met
- Qual samples delivered to AI data-center customers (100s per DDR4 precedent)

ASSUMPTIONS

- Foundry Partner 1B process is frozen; no major PDK or process changes expected after fix
- Tape-out split into base die + memory die; ≤2 all-layer iterations per die before WS retrofit
- ALD process flow via ALD Vendor at Research Lab (Path A struck per 21 Apr decision)
- Base die: Client Co designs; Logic Foundry implements; Interface Design Partner digital implementation support
- 4H W2W hybrid bonding with TSVs is viable at 16 Gb per memory layer (Foundry Partner confirms)
- Foundry Partner packages memory with Logic Foundry/Interface Design Partner base die; Backup Foundry is contingency path only
- Stacking partner supplies ADK to avoid manual stitching; face orientation decided
- Adequate testing capability available; test hardware to opex (no capex for testers)
- FS → WS retrofit pattern valid; 2nd Si incorporates FS learnings without re-architecture
- ELFR study for 20 dppm follows MP release; post-qual, not gating
- Fab cycle time <2 months for full flow on a hot lot

RISKS

- **D5-01** Ferroelectric cell fails product KPIs — Cell Lead
- **D5-02** Foundry Partner PDK non-functional in EDA Vendor B flow — blocks design start — Design Lead
- **D5-03** Product spec not closed before design start — already slipped — Spec Owner
- **D5-04** Hybrid-bonding technology not selected — blocks partner, stacking, base-die tech — Fab/Process Lead
- **D5-05** Base die and memory die asynchronous — 'two-bridge' sync risk — Product Boss
- **D5-15** 4H W2W hybrid bonding yield/reliability at production density is unknown
- **D5-20** Qual failure on embedded flow requiring design/process changes
- **Geo** Foundry Country concentration (Foundry Partner + Logic Foundry + Interface Design Partner + stacking) — geopolitical exposure
- **Cost** Program cost envelope not yet defined; no baseline for quarterly review
- **U5** DDR5 IP source — vendor and license model open — Design Lead

NEXT STEP · Foundry Country trip with Foundry Partner is the fulcrum — closes hybrid bonding tech, base-die partner, and PDK escalation. Rev 2.0 of this charter with locked calendar dates follows in early May 2026.

Milestone detail

Doneness criteria, key subtasks, assumptions, and linked risks per milestone — M1 through M5 (Rev 1.4 schedule).

M1 Jun 2027	M2 Mar 2028	M3 Dec 2028	M4 Jun 2029	M5 Dec 2029
<p>Ferroelectric cell validated</p> <p>OWNER Cell Lead (Process & Integration)</p> <p>EFFORT SL vehicle + process development at Research Lab</p> <p>DONENESS CRITERIA</p> <ul style="list-style-type: none"> Preliminary specifications signed off Short-loop (SL) vehicle developed at Research Lab Integration path determined (ALD Vendor via Research Lab) Reduced cap height integrated Low-temperature ALD process developed Ferroelectric cell meets product KPIs <p>KEY SUBTASKS</p> <ul style="list-style-type: none"> Execute SL vehicle tapeout; run learning cycles Monthly cell-KPI review from May 2026 Integration DOE for reduced cap height ALD recipe development + characterization <p>ASSUMPTIONS</p> <ul style="list-style-type: none"> ALD Vendor JDP in place; IP ownership resolved Research Lab SL vehicle operational by Q2 2026 Ferroelectric material passes reliability screen <p>LINKED RISKS D5-01, D5-03, D5-13, D5-14</p>	<p>FS functional 3D stack validated</p> <p>OWNER Design Lead (Design) Fab/Process Lead (Fab/Process)</p> <p>EFFORT Compound milestone — 4 sub-gates</p> <p>DONENESS CRITERIA</p> <ul style="list-style-type: none"> Base die delivered: technology, foundry, interface provider selected; base die taped out, sorted, tested Memory die delivered: Foundry Partner PDK received; memory die taped out, sorted, tested 1st Si 3D stack delivered: hybrid bonding selected; stacking partner chosen; stack sorted + tested 1st Si 3D stack validated: parametric + functional tests pass; DDR5 at 3200 MT/s; TSV verified; retention met <p>KEY SUBTASKS</p> <ul style="list-style-type: none"> Foundry Country trip closes hybrid-bonding + base-die partner Base die tapeout to Logic Foundry via Interface Design Partner Memory die tapeout to Foundry Partner 1B W2W hybrid bonding at stacking partner; ADK integration Sort + test programs developed for base, memory, stack <p>ASSUMPTIONS</p> <ul style="list-style-type: none"> Foundry Country trip closes hybrid-bonding technology decision Logic Foundry + Interface Design Partner signed; PDK available for base die 1st Si returns within 12 weeks of tapeout Stack yield sufficient for parametric characterization <p>LINKED RISKS D5-02, D5-04, D5-05, D5-07, D5-08, D5-09, D5-10, D5-11, D5-12</p>	<p>WS functional 3D stack validated</p> <p>OWNER Design Lead (Design) Fab/Process Lead (Fab/Process)</p> <p>EFFORT 240 man-months retrofit + 2nd Si ramp to 10 eng</p> <p>DONENESS CRITERIA</p> <ul style="list-style-type: none"> Retrofit base + memory tapeouts incorporate FS learnings 2nd silicon validated through full DDR5 protocol Stack-level variation budget confirmed across 4H CXL + DIMM interface pre-validation complete <p>KEY SUBTASKS</p> <ul style="list-style-type: none"> FS learning distillation → retrofit spec freeze 2nd Si tapeout with incremental changes only Full DDR5 protocol characterization at 3200 MT/s Stack-level SPC rollout <p>ASSUMPTIONS</p> <ul style="list-style-type: none"> FS learnings are incremental, not architectural Retrofit scope ≤1 FE change + ≤2 metal changes per die Fab cycle time <2 months on hot lot No capex required for testers <p>LINKED RISKS D5-06, D5-15, D5-18, D5-25</p>	<p>ES functional 3D stack validated</p> <p>OWNER Product Lead (Product) Design Lead (Design)</p> <p>EFFORT 18 man-months sample build + validation</p> <p>DONENESS CRITERIA</p> <ul style="list-style-type: none"> Engineering samples built on production-intent flow Full JEDEC DDR5 compliance demonstrated CXL and DIMM interfaces validated at system level Samples shipped to AI data-center customers <p>KEY SUBTASKS</p> <ul style="list-style-type: none"> Customer sample quantity finalized; build plan locked Production-intent flow frozen at WS boundary Customer-facing datasheet + qual report drafted Engineering sample ship approval (LAR or equivalent) <p>ASSUMPTIONS</p> <ul style="list-style-type: none"> Customer sample quantity ~100s (DDR4 precedent) Production-intent flow frozen at WS milestone No respin required between WS and ES <p>LINKED RISKS D5-08, D5-15, D5-20, D5-25</p>	<p>Product qualified — release</p> <p>OWNER Product Boss</p> <p>EFFORT 18 man-months qualification + release</p> <p>DONENESS CRITERIA</p> <ul style="list-style-type: none"> JEDEC qual passed on production flow Retention 3 months @ 85°C verified Endurance 10^15 cycles confirmed UBER 10^24 achieved Released to production; datasheets published <p>KEY SUBTASKS</p> <ul style="list-style-type: none"> Qual test program execution per RCM Reliability stresses completed; delta analysis signed Characterization test program completed Production stocking plan in place H-toll 1000 h weighting review concluded <p>ASSUMPTIONS</p> <ul style="list-style-type: none"> Qual samples delivered to customer on schedule No respin required post-ES ELFR study for 20 dppm follows release, not gating <p>LINKED RISKS D5-09, D5-11, D5-21, D5-22</p>

Incremental project breakdown

Parsing the DDR5 NV-DRAM+ program into a roadmap of narrower-scope projects with aggressive intermediate deliverables.

FRAMING — the case for incremental scope

Boil-the-ocean programs fail. The 21 April milestone scrub already flagged this: 3D stack and monolithic 16 Gb should be separate projects, and the base die and memory die could be separate projects until bonding. The current 4H full-product plan runs ~4 years with five compound milestones, each gated by a prior unknown. Breaking the program into four narrower tracks retires risks sequentially, produces shippable or licensable deliverables at every gate, and creates forcing functions on decisions that currently have no deadline. Aggressive interim dates replace the single 2029 commitment with a cadence of customer-visible outcomes.



TRACK A	TRACK B	TRACK C	TRACK D
Q3 2027	Q1 2028	Q3 2028	Q4 2029
Ferroelectric cell + 16 Gb monolithic DDR5 NV-DRAM <p>SCOPE Validate ferroelectric cell in Foundry Partner 1B; ship a monolithic (non-stacked) 16 Gb DDR5 NV-DRAM product.</p> <p>DELIVERABLE Commercial 16 Gb DDR5 NV-DRAM with 3-month retention; characterization report; release-to-production package.</p> <p>RISKS RETIRED</p> <ul style="list-style-type: none"> Ferroelectric cell KPI risk (D5-01) Foundry Partner 1B process-flow integration risk Product-spec closure gap (D5-03) Initial DDR5 IP license decision (U5) <p>WHY THIS STEP No 3D integration, no hybrid bonding, no base die required. Leverages the cell validation already in flight at Research Lab. Gives Client Co a commercial DDR5 NV-DRAM presence 18 months before the 4H product.</p> <p>TIME TO DELIVERY Q3 2027 · ~18 months from charter</p> <p>DEPENDS ON M1 cell validation; Foundry Partner PDK fix (U1)</p>	Base die platform qualification <p>SCOPE Qualify the Client Co-designed base die at Logic Foundry via Interface Design Partner as a standalone IP block, decoupled from the memory die program.</p> <p>DELIVERABLE Qualified base die PDK + characterization report; interface specification frozen; base die as reusable block across density variants.</p> <p>RISKS RETIRED</p> <ul style="list-style-type: none"> Base-die partner/technology risk (D5-04) Asynchronous two-bridge risk (D5-05) Backup Foundry-vs-Logic Foundry-vs-other node decision (U7) ADK availability from stacking partner (U8) <p>WHY THIS STEP The base die is common across all density variants. Qualify once, use many times. Decouples base-die schedule from memory die so Track A can proceed even if base die slips. Enables IP licensing as a side revenue line.</p> <p>TIME TO DELIVERY Q1 2028 · ~21 months from charter</p> <p>DEPENDS ON Foundry Country trip decisions; Logic Foundry + Interface Design Partner partnership signed</p>	2H hybrid-bonding pilot (32 Gb stacked) <p>SCOPE Build and validate a 2-high stack using the qualified base die + a 16 Gb memory die over W2W hybrid bonding.</p> <p>DELIVERABLE 32 Gb DDR5 NV-DRAM+ (2H) on production-intent hybrid-bonding flow; full JEDEC compliance at 2H density; shipping mid-density SKU.</p> <p>RISKS RETIRED</p> <ul style="list-style-type: none"> Hybrid-bonding yield at low stack height (D5-15) Stack-level variation budget across 2 layers TSV connectivity + retention at stack level Sort + test program scale-up for stacks <p>WHY THIS STEP 2H stack halves the integration risk of 4H and produces a shippable mid-density product. Proves the bonding flow under a less stressed yield envelope. Customers targeting mid-density get a product without waiting for 4H.</p> <p>TIME TO DELIVERY Q3 2028 · ~29 months from charter</p> <p>DEPENDS ON Track A (memory die proven); Track B (base die qualified)</p>	4H 64 Gb DDR5 NV-DRAM+ — original program <p>SCOPE Integrate the validated cell, qualified base die, and proven bonding flow into the full-density 4H 64 Gb product.</p> <p>DELIVERABLE JEDEC-qualified 64 Gb DDR5 NV-DRAM+ (4H) for AI data-center CXL + DIMM workloads, at 3-month retention and 10¹⁵ endurance.</p> <p>RISKS RETIRED</p> <ul style="list-style-type: none"> 4H stack-level variation budget Reliability at production density Manufacturing flow at target yield Full customer qualification <p>WHY THIS STEP Integration target. By this point each prior track has retired a specific unknown; 4H is the final assembly of proven components rather than a single giant unknown. Same customer commitment window as the original plan — but with three shipping products behind it.</p> <p>TIME TO DELIVERY Q4 2029 · ~48 months from charter</p> <p>DEPENDS ON Tracks A, B, C all closed</p>

BENEFITS OF INCREMENTAL ROADMAP

- Each track produces a shippable or licensable deliverable — business value at every gate
- Risks retire sequentially, not simultaneously — no compound failure mode
- Early revenue from Tracks A + C offsets Track D program cost
- If Track D slips or fails, Tracks A–C still ship — business does not go to zero
- Aggressive interim dates force decisions that today have no deadline

DECISIONS REQUIRED TO ADOPT

- Adopt incremental roadmap, or continue boil-the-ocean 4H-only approach
- Approve split: 16 Gb monolithic and base-die platform as separate tracks
- Staffing: sequential (single team) vs. parallel (separate teams per track)
- Budget authorization — gate-by-gate increments vs. full program up front
- Customer communication: whether to announce interim products before 4H

Design team sizing & phasing — industry benchmarks

Peak headcount, schedule duration, and SME composition for DDR memory chip design — with application to the Foundry Partner 1B 16Gb 3D DDR5 NV-DRAM+ program.

16 Gb DDR on ~19 nm — 1y-nm class, DDR5 era

Comparable to Foundry Partner 1B 16Gb 3D (monolithic die portion)

PEAK DESIGN TEAM

90 – 110 engineers

ARCHITECTURE → FIRST SI

30 – 36 months

FIRST SI → QUALIFIED

12 – 18 months

Complexity drivers: DDR5 protocol (mode registers, sub-channels), DFE I/O, on-die ECC, DQ speeds ≥ 6400 MT/s, tight retention/refresh budgets at advanced node. Industry references: Tier-1 Memory Vendors 1y-nm 16 Gb designs.

TOTAL PROGRAM

42 – 54 months

ENGINEER-MONTHS (TOTAL EFFORT)

3,000 – 3,800

EXPECTED TAPE-OUTS

2 all-layer + 1–2 metal

4 Gb DDR on 25 nm — DDR3 / early-DDR4 era

Comparable to the Client Co sister program (Sister Foundry 25 nm 4 Gb DDR4 NV-DRAM)

PEAK DESIGN TEAM

30 – 45 engineers

ARCHITECTURE → FIRST SI

20 – 26 months

FIRST SI → QUALIFIED

9 – 12 months

Note: 25 nm is below DDR5's typical density floor (JEDEC DDR5 starts at 8 Gb); benchmark is most relevant to DDR3 / early-DDR4 generation. ~3–4x smaller effort vs the 16 Gb / 19 nm case: 4x density, simpler I/O, simpler protocol state.

TOTAL PROGRAM

29 – 38 months


ENGINEER-MONTHS (TOTAL EFFORT)

800 – 1,100

EXPECTED TAPE-OUTS

1 all-layer + 1 metal

SME composition of a ~100-engineer DDR5 design team — typical discipline allocation at program peak



Physical design / layout

Floorplan, P&R, DRC/LVS, parasitic extraction, power grid, clock tree synthesis

Periphery / control design

Command decode, control logic, mode registers, redundancy, DFT scan, fuse logic

Array / core design

Memory cell array, sense amps, word/bit-line drivers, row/col decoders

Verification (FV + AMS)

UVM testbenches, JEDEC compliance, mixed-signal simulation, STA, reliability sign-off

High-speed I/O

DQ drivers / receivers, DFE, DLL/PLL, ZQ calibration, CA I/O

Analog / mixed-signal

VPP / VDD generators, LDOs, bandgap, references, sense-amp bias


Architect / PM

Spec ownership, block partitioning, IP selection, schedule, risk management

Product / test engineering

ATE pattern gen, DFT insertion, silicon bring-up and debug support

Time-phased staffing curve — engineer headcount by program month, with phase bands



APPLIED TO OUR PROGRAM — implications for Foundry Partner 1B 16 Gb 3D DDR5 NV-DRAM+ staffing

- **Benchmark peak:** ~90–110 design engineers for the monolithic 16 Gb die alone. 3D stacking, base die, and hybrid-bonding workstreams add ~15–25 more at integration peak.
- **Schedule fit:** Rev 1.4 baseline (Q1 2026 start → M5 Dec 2029 = 48 mo) now aligns with the 42–54 mo benchmark range for 16 Gb 1y-nm — realistic under the assumption of zero-respin WS→ES and a base die reused from Track B.
- **Current gap:** the named core team is 8 people (Product Boss, TPM, Design Lead, Fab/Process Lead, Device Lead, Post-Fab Lead, Product Lead, Exec Sponsor). The delta to a 90–110-engineer peak is the real staffing plan — not yet sized in the charter.
- **Incremental-track implication (page 3):** Track A (16 Gb monolithic) requires ~40–60 engineers peak — a 4 Gb / 25 nm team size will NOT cover a 16 Gb / 19 nm monolithic product, even with density reduced.

Budget estimate — program development cost

Wafer development, FTE (Region B + Asia split), and consulting for the 48-month Foundry Partner 1B 16 Gb 3D DDR5 NV-DRAM+ program.

BASELINE TOTAL — 48-MONTH PROGRAM

Range: \$60M – \$97M

~\$75M

Rev 1.4 schedule: Q1 2026 → M5 Dec 2029

EXCLUDED FROM THIS ESTIMATE

Process & metrology tools at foundry partner line (TBD)
 Fab buildout, GTM / launch, post-MP sustaining engineering
 — see detailed exclusions at bottom of page

Wafer development

Short-loop vehicle + integration (Research Lab, ~300 wafers)	\$6.0M
Mechanical / process qual wafers	\$1.5M
Base die tape-outs (3 spins @ Logic Foundry via Interface Design Partner)	\$4.0M
Memory die tape-outs (3 spins @ Foundry Partner 1B)	\$7.0M
Protolots (base + memory, POR flow)	\$3.0M
Hybrid-bonding test lots (W2W stacking partner)	\$5.0M
Engineering-sample + JEDEC qual lot wafers	\$3.5M

Subtotal **\$30M**

Range \$21M (FTR learning curve compresses, zero respins) to \$45M (stacking respins, extra qual lots). Assumes 2-3 all-layer + 1-2 metal tape-outs per die.

FTE — engineering

Total engineer-months (48 mo program)	~3,400
REGION B — ~35% share (~1,190 em)	
Architecture, device, integration, process, PM	
Fully-loaded rate: ~\$17K / month	\$20.2M
ASIA — ~65% share (~2,210 em)	
Physical design, verif, layout, periphery, test	
Blended Foundry Country / Region D / Region E rate: ~\$9K / month	\$19.9M

Subtotal **\$40M**

Range \$35M (if Asia share grows to 75%, or engineer-months 15% lower) to \$45M (if peak runs longer or Region B share grows). Fully-loaded = salary + benefits + overhead + facilities.


Consulting

DDR5 IP licensing + vendor support	\$2.0M
Hybrid-bonding external expertise	\$0.8M
EDA tool-flow consulting (advanced node)	\$0.7M
JEDEC compliance / verification consulting	\$0.4M
Reliability / qual labs (external test)	\$0.6M
Program advisory (Strategic Advisor et al.)	\$0.5M

Subtotal **\$5M**

Range \$4M to \$7M depending on whether DDR5 PHY/controller is licensed vs in-house, and scope of external reliability testing. Program advisory includes Strategic Advisor TPM and methodology.

Time-phased spend — quarterly burn rate by category, with cumulative trajectory



EXCLUDED FROM THIS ESTIMATE

- Process and metrology tools added to the foundry partner line (Foundry Partner, Logic Foundry, stacking partner) — TBD; driven by decisions at Foundry Country trip and subsequent scrub.
- Facilities / fab buildout at partner sites; clean-room capex.
- GTM, marketing, launch activities, channel and customer ramp support.
- Post-MP sustaining engineering, ELFR study execution, second-source development.
- Royalties on third-party IP (amortized over unit shipments, not booked as R&D).

KEY ASSUMPTIONS & CONFIDENCE

- Currency: USD throughout; baseline 2026 prices; no inflation escalator applied.
- Geographic split: 35% Region B / 65% Asia, weighted toward Region B for architecture and process, Asia for layout and verification.
- Fully-loaded rates include salary, benefits, overhead, and facilities at each location.
- Confidence: industry-benchmark range, not a Client Co-costed plan. Rev 1.1 of this page must validate against internal finance model before commitment.

Inputs required to start — project-start prerequisites

Everything that must be in place on day one — technology data, agreements, team, spec, and governance — before the program can credibly begin.

REV 1.1 · 24 APRIL 2026

Updated schedule — realistic baseline

DEFINITION OF READY — the distinction between assumptions and prerequisites

Assumptions (page 1) are stated beliefs that may or may not hold. Prerequisites are deliverables that must physically exist and be signed off. The program should not advance past the Q1 2026 start gate until every item below is either in hand (green), with a committed date (amber), or has an explicit mitigation owner (red). Items flagged red are the short list of what must close before PO release to Foundry Partner, before first wafer out, and before the M1 cell validation milestone.

READINESS LEGEND:

● Ready / in hand

● Partial / committed date

● Gap / mitigation required

Status reflects assessment as of 24 April 2026; to be re-scrubbed weekly and at each phase gate.

Technology — from Foundry Partner

- **PDK — production-ready**
Digital + AMS views, major-vendor EDA flows. Currently non-functional in primary EDA flow — blocks design start. D5-02 / U1.
- **Design Rule Manual (DRM)**
Latest revision with 3D-stack annotations; changes from 1B early draft must be delta-documented.
- **SPICE models + corners**
TT / FF / SS / SF / FS + Monte Carlo; aging (BT1, HCI); temperature 0–85°C. Foundry Partner to deliver at PDK v1.0.
- **Tech files (LEF, Milkyway, techLEF)**
For place-and-route flow. Expected with PDK v1.0.
- **Parasitic extraction decks**
QRC / StarRC tech files at multiple corners. With PDK.
- **DRC / LVS sign-off decks**
Sign-off decks; interim decks acceptable for block-level work.
- **WAT + SPC coverage**
Wafer acceptance test spec, statistical process control data — needed to define redundancy + repair strategy.
- **Fuse / redundancy spec**
Row / column repair scheme + fuse map definition. Foundry Partner to deliver Q3 2026.

Technology — partners

- **ALD Vendor ferroelectric ALD process**
Specs, integration flow, chamber qualification status. Path via Research Lab struck for alternative ALD per 21 Apr decision.
- **Research Lab SL vehicle data**
Short-loop integration results, yield baseline, process window — needed to establish cell KPIs.
- **Logic Foundry (via Interface Design Partner) base die PDK**
Base die process at Logic Foundry on 12LP+ equivalent; design rules, IP library. Partnership not yet signed.
- **Stacking partner ADK**
Assembly Design Kit from W2W hybrid-bonding partner; face orientation + TSV rules. Partner TBD pending Foundry Country trip.
- **Hybrid-bonding design rules**
Bond-pad pitch, alignment tolerance, overlay budget, retention tests. Partner-specific; TBD.
- **EDA tool versions**
Major EDA vendor toolchains for layout, sign-off, and verification — versions compatible with Foundry Partner 1B PDK.
- **JEDEC DDR5 spec version**
Target JESD79-5B or newer; scope of optional features (DFE taps, on-die ECC) locked.

Legal + commercial agreements

- **JDA: Client Co ↔ Foundry Partner**
Joint Development Agreement — scope, IP ownership, milestones, cost share, exclusivity window. Draft circulating.
- **JDP: Client Co ↔ ALD Vendor**
Ferroelectric ALD process development. IP ownership resolution needed. D5-13.
- **Service contract: Research Lab**
Short-loop vehicle development, integration support. Statement of Work to be finalized.
- **Tech transfer + services: Logic Foundry + Interface Design Partner**
Base die design services from Interface Design Partner, manufacturing at Logic Foundry. Terms + pricing open.
- **MOU / JDA: stacking partner**
W2W hybrid bonding. Partner identity, terms, and IP position all open pending Foundry Country trip.
- **NDA coverage — all parties**
Bilateral NDAs in place for Foundry Partner, ALD Vendor, Research Lab, Interface Design Partner, Logic Foundry; multi-party NDA for stacking partner once identified.
- **DDR5 IP license**
PHY + controller IP source (vendor vs in-house). License model open. U5.
- **Export controls (EAR, EU dual-use)**
Jurisdictional review for memory + bonding tech; clearances where required.

Team · infra · spec

- **Full staffing plan**
Named core team is 8; benchmark peak is ~100–110. Staffing plan to close the delta is not yet written. See page 4.
- **EDA licenses at peak scale**
License pool sized for ~100-engineer peak across major EDA vendors. Cost in consulting bucket.
- **Compute cluster**
Full-chip P&R, STA, SPICE simulation capacity. Existing cluster may need expansion at Y2 peak.
- **Physical site security**
Export-controlled design bays, clean-room access at partner sites. Reviewed at program start.
- **Product spec frozen**
Product spec not closed before design start — already slipped. Owner: Spec Owner. D5-03.
- **Customer target list + volume forecast**
AI data-center customers + initial volume assumptions; feeds business case + capacity planning.
- **Business case / ROI model**
Charter covers cost side; revenue side + margin model needed before full commitment.

Funding + governance

- **Approved program budget**
Baseline ~\$75M (range \$60–97M) per page 5. Rev 1.1 of budget page must validate against Client Co finance model before commitment.
- **Process / metrology tool capex**
Foundry-partner line tooling additions — TBD; excluded from current budget estimate. Must be sized post Foundry Country trip.
- **Phase-gate governance**
SWIFT-style gate structure defined (Gate 2/3/4/5/6); gate criteria and voting membership to be baselined.
- **Exec steering committee**
Chartered under Exec Sponsor. Monthly cadence, mandate and escalation path documented.
- **Risk register baselined**
25 tracked risks (9 Critical, 11 High) in DDR5 register as of 21 Apr. Weekly review cadence in place.
- **Quarterly budget review**
Q1 2026 burn vs plan review cadence; variance thresholds + re-forecast triggers defined.
- **Foundry Country trip + Foundry Partner scrub**
Planned May 2026. Fulcrum decision gate — closes hybrid-bonding tech, base-die partner, and PDK escalation.

THE SHORT LIST — red-status items that must close before project can be declared “READY TO START”

- Foundry Partner PDK operational in primary EDA flow + Design Rule Manual released — blocks design start. D5-02.
- Hybrid-bonding partner selected with ADK + design rules in hand — blocks 3D stack architecture, base-die partner, stacking flow. D5-04. Foundry Country trip (May 2026) is the forcing function.
- Base die partnership signed (Logic Foundry + Interface Design Partner terms + pricing) — blocks Track B and full-program M2 tapeouts.
- Product spec frozen (owner: Spec Owner) — already slipped per 21 April scrub. D5-03.
- Budget approval + exec sponsor mandate — PO authority to commit to Foundry Partner, Research Lab, ALD Vendor, and Interface Design Partner service contracts.

Inputs mapped to schedule — when each prerequisite must be in hand


Every item from page 6 placed at the quarter it must close to avoid becoming a critical-path blocker. Color = current readiness (green/amber/red from page 6).

HOW TO READ

Each row in a category lane is one prerequisite from page 6. Each marker sits in the quarter by which the item must be signed off and in hand for downstream work to start on time. A red marker means the item is currently a gap and its need-by date is approaching. The gate-load strip at the top counts prerequisites due per quarter — Q1 2026 carries the heaviest load because most items must close at project start. Vertical dashed lines mark M1–M5 milestones and the May 2026 Foundry Country trip.

GATE LOAD

items due / quarter



KEY DROP-DEAD DATES

Q1 2026 18 items due

Project-start gate — PDK, DRM, all core agreements, spec freeze, budget, governance

Q2 2026 Foundry Country trip + 6 items

Hybrid-bonding partner selection, ALD process, SPICE models, WAT coverage, volume forecast

Q4 2026 Base die partnership gate

Logic Foundry + Interface Design Partner terms signed, base die PDK available, parasitic extraction decks

Q2 2027 Stacking partner gate

ADK + hybrid-bonding rules in hand before M1 cell validation (Jun 2027) and base-die tapeouts

LEGEND: ◆ Ready / in hand ◆ Partial / committed date ◆ Gap / mitigation required Gate-load heat (quarterly count, low → high)

Use this page to drive a weekly pre-start scrub: walk from Q1 2026 forward; for each red marker, either upgrade status or re-baseline its need-by date. Do not accept "item pushed right" without checking the downstream lane it feeds.

References & methodology

Sources behind the resourcing and budget estimates — pages 4 and 5. Industry-benchmark synthesis, not a Client Co–costed plan.

REV 1.1 · 24 APRIL 2026

Methodology & sources — added in Rev 1.1

ESTIMATION APPROACH — three principles that govern every figure on pages 4 and 5

1. Range, not point.

Every figure carries a low/high band. The point estimate is the midpoint of converging sources, not a single vendor quote.

2. Two-source minimum.

No headcount, rate, or unit cost is used unless at least two independent sources agree within ±25%. Where sources disagree, the wider range is shown.

3. Public benchmark, not vendor RFQ.

These figures come from public industry research and peer-program benchmarks — not Foundry Partner, Logic Foundry, Interface Design Partner, or stacking-partner quotes. Vendor quotes replace them in Rev 2.0.

Methodology — how each estimate was built

How resourcing (page 4) was derived

- Peak headcount range pulled from published 16 Gb / 1y-nm DDR5 design programs at Tier-1 memory vendors — disclosed indirectly via author counts on industry conference papers, technology disclosure sessions, and analyst headcount tracking.
- SME mix percentages (22% physical / 20% periphery / 18% array / 12% verif / 12% I/O / 8% AMS / 5% architect / 3% test) are a composite from publicly indexed SoC and DRAM team allocations, cross-checked with two recent peer programs.
- Engineer-months are headcount × duration, not bottom-up from a work breakdown. Bottom-up sizing arrives in Rev 2.0 from the staffing-plan owner.
- Time-phased curve (Architecture → Block → Integration → T/O → Si debug → Qual) follows the standard chip-program shape; phase boundaries calibrated to Rev 1.4 milestones.

How budget (page 5) was derived

- Tape-out cost = mask-set cost (foundry-published ranges for advanced nodes) + first-lot wafer cost. Assumes 2–3 all-layer + 1–2 metal spins per die.
- FTE rates are fully-loaded (salary + benefits + overhead + facilities), regional 2026 prices, no inflation escalator. Region B ~\$17K/mo and Asia blended ~\$9K/mo are midpoints of three independent regional surveys.
- Hybrid-bonding lot cost assumes industry-typical W2W pricing at production-density pitch — the largest single uncertainty in the wafer line.
- Consulting line items are scoped at the work-package level, not by day rate. DDR5 IP licensing dominates.

Confidence per category

- Wafer development: ±25% — mask-cost variance is the largest driver.
- FTE engineering: ±15% — rates are well-benchmarked; total engineer-month count drives the variance.
- Consulting: ±30% — depends on IP licensing path (vendor vs in-house) and external lab scope.

Resourcing benchmarks — page 4 sources

Standards & specifications

- Industry standard DDR5 SDRAM specification (defines protocol complexity that drives team size).
- Industry standard HBM3 specification (hybrid-bonding precedent for stack discipline).
- Industry 3D-IC integration standards series.

Industry research reports

- Cost-of-IC-design research firm — advanced-node report series.
- Node-level economics analyst — IC cost model and node-level economics, recent editions.
- 3D integration & hybrid-bonding research firm — memory industry & market trends; annual reports.
- Teardown analyst — DRAM teardown analyses (1y-nm class, 16 Gb DDR5 SKUs).
- Memory R&D analyst — published estimates on memory program R&D scale.

Conference proceedings

- Industry conference proceedings — DDR5 sessions; 1α / 1β / 1γ-class papers from Tier-1 memory vendors.
- Symposium on technology & circuits — process and circuit papers, same vendors.
- Electron device meeting — 1y-nm DRAM cell scaling and integration papers.
- Design & test journal — chip-program productivity studies.

EDA / IP product data

- IP Vendor A DDR5 PHY + controller — public product briefs.
- IP Vendor B DDR5 PHY + memory controller IP — public product briefs.
- IP Vendor C DDR5 PHY — public datasheets.

Budget benchmarks — page 5 sources

Mask & tape-out cost

- Cost-of-IC-design research firm — mask-cost-by-node tables, 1y-nm class, current editions.
- Node-level economics analyst — wafer cost, mask cost, lot pricing models.
- Foundry public guidance: Tier-1 foundries annual reports + investor updates; Logic Foundry base-die node cost via Interface Design Partner channel.
- Memory R&D analyst — published mask + wafer cost estimates by node.

Hybrid-bonding & 3D integration

- 3D integration & hybrid-bonding research firm — cost-of-ownership reports.
- Industry 3D IC working group — pricing benchmarks at production density.
- Hybrid-bonding tool & IP vendors — public statements on hybrid-bonding cost trajectories and tooling.

Engineer rates by region

- Region B: regional compensation surveys (3 independent sources).
- Foundry Country: semiconductor salary surveys + market intelligence.
- Region D: semiconductor compensation surveys + tier-1 recruiter data.
- Region E: industry strategic reviews + tier-1 MNC published bands.

DDR5 IP licensing & external services

- IP Vendors A / B / C — public IP product pricing where disclosed; analyst commentary on DDR-class license cost ranges.
- Reliability / qual labs: industry standard day-rate bands for JEDEC qual programs; ELFR study cost benchmarks from peer DRAM programs.

Internal & program references

Internal Client Co documents, Strategic Advisor methodology, and the Rev-1.4 schedule baseline that drive the duration assumptions on pages 4 and 5. These are the program-specific inputs that calibrate the public benchmarks.

Client Co program documents

- DDR5 product brief — Persistent Memory Solutions for Efficient AI Scale-out (Client Co).
- Go-to-Market Core Teams — Foundry Partner 1B 16 Gb 3D DDR5 NV-DRAM track (Exec Sponsor; Product Boss; TPM; et al.).
- DDR5 Risk Register — 25 risks (9 Critical, 11 High), 21 April 2026 baseline.
- 21 April 2026 milestone scrub — source of the Rev 1.4 schedule that drives p4 and p5 duration assumptions.
- DDR5 estimates workbook (Client Co internal).
- Client Co sister program — Sister Foundry 25 nm 4 Gb DDR4 NV-DRAM: peak headcount, spend, and schedule comparator.
- Foundry Partner PDK status disclosures — primary EDA flow non-functional issue, D5-02.
- ALD Vendor JDP draft — alternative ALD path, 21 April decision.

Strategic Advisor methodology

- FTTM (Fast Time To Market) framework — phase-gate structure, milestone discipline.
- Strategic Advisor project charter standard — the structure of this document.
- Strategic Advisor SOW templates and white-paper standard.

Charter prior revisions

- Rev 1.0 — initial baseline (early April 2026).
- Rev 1.1 (this revision) — 24 April 2026, updated schedule realistic baseline; references page added.
- Rev 1.4 schedule — referenced in p2 and elsewhere as the committed milestone calendar.
- Rev 2.0 — planned post Foundry Country trip (May 2026); replaces benchmark-based costs with vendor quotes.

RECONCILIATION PLAN — what must close before Rev 2.0 commits the numbers

1. Vendor quotes replace benchmark ranges.

Wafer line items switch from public-benchmark midpoint to actual quotes from Foundry Partner, Research Lab, ALD Vendor, Logic Foundry/Interface Design Partner, and the stacking partner. Expected delta: ±20% per line.

4. IP licensing decision.

DDR5 PHY + controller: vendor (IP Vendor A / B / C) vs in-house. Outcome moves the consulting line by \$1–3M.

2. HR rate validation.

Client Co HR provides actual fully-loaded rates by location and seniority — replaces the ~\$17K / \$9K regional midpoints. Expected delta: ±10%.

5. Process / metrology tool capex.

Currently excluded from p5; sized post Foundry Country trip. Could shift the total envelope by 10–25% if treated as program cost rather than partner capex.

3. Bottom-up staffing plan.

Delta from 8-person core to ~110-engineer peak is a real plan owned by TPM and Product Boss — not an industry curve overlay. Engineer-months recompute against that plan.

6. Inflation escalator.

2026 prices held flat across 4 years; finance to apply standard escalator in Rev 2.0 — likely +3–6% over the program window.

GOAL

Release Driver Node X driver with both Legacy FET Fab and Target FET Fab FETs Project Beta by Q4'2027

SUCCESS CRITERIA

9 items

- Gen 3 SGMOS FETs successfully ported from 8-inch Legacy FET Fab to 12-inch Target FET Fab fab with process qualified by Q1 2027
- Standalone Driver Node X BCD Lite HPC driver IC reaches final silicon tape out by Q2 2027 within ~2 all-layer tapeouts (A0, B0)
- Performance equivalence demonstrated between Project Alpha and Project Beta modules
- 12-inch end-to-end module manufacturing flow proven out including FSM, BG/BSM (glass bonding), probe, RDL+bump, embedding, and assembly
- Full JEDEC and End Customer-specific qualification passed (not by similarity to Project Alpha)
- PCN approval obtained from Customer Tier customers prior to MP release (Prior or after?)
- Target gross margin of 50% achieved through internalization and COGS reduction
- Target FET Fab fab able to autonomously support production volumes with all tooling gaps closed
- ELFR study confirms 20 dppm target on production candidate module (post-MP)

ASSUMPTIONS

9 items

- SGMOS Gen 3 FET Legacy FET Fab-to-Target FET Fab port: process qual target Q1 2027. Requires CAPEX approval and tooling investment
- Driver Node X BCD HPC Lite foundry development is ongoing and will not affect critical path - driver silicon expected ready before FET fab readiness
- Standalone Driver IC for Project Beta requires ~2 all-layer tapeouts (A0, B0). PDK version and T/O date need to be confirmed
- Initial Driver Node X tapeout planned for August. Preliminary build will use Project Alpha FETs (Legacy FET Fab)
- Target FET Fab will be ready to start running protolots by Q1 2027.
- BG/BSM for 12-inch will use glass bonding (differs from Project Alpha Taiko-based process)
- 12-inch can do probe at Mid-end OSAT. RDL+Bump at Bump OSAT
- POD is FET wafer size agnostic - 8-inch FET + 12-inch driver lookahead build is viable
- Qual and validation must be repeated due to mfg flow differences - cannot qualify by similarity from Project Alpha

RISKS

8 items

- Target FET Fab process qualification delay - any slip from Q1 2027 target pushes the entire Project Beta timeline since Target FET Fab FETs are the gating dependency
- Driver Node X PDK maturity - foundry must deliver PDK at sufficient maturity for design tapeout; delays here impact driver silicon schedule
- Gen 3 FETs may not meet expected metrics on 12-inch Target FET Fab, particularly if 50-micron backgrind is required. Gen 2.5 FETs are fallback
- CAPEX approval and tooling investment for Target FET Fab - required for production volume support, timeline depends on EC approval
- Glass bonding process for 12-inch BG/BSM is new and differs from Project Alpha Taiko process - yield and reliability risk
- TDC thermal performance - risk of lower thermal dissipation requiring heat slug redesign or mold compound changes
- Mold void under IC driver - requires DOE for mold parameters and mold flow simulation at assembly sites
- Embedded substrate product for NPI is new - high reliability risk requiring look-ahead reliability data from Reference Program A/Reference Program B/Project Alpha programs

M1 12-inch Module Flow

Mid-End Capability for 12 inch (post fab processing FET- FSM, BGBM , Driver - RDL/Bump)

DONENESS (5)

- 8/12-inch Target FET Fab FSM capability established at Mid-end OSAT
- 8/12-inch Target FET Fab BG/BM with glass bonding proven out at Mid-end OSAT
- 12-inch probe hardware verified for FET
- 12 -inch probe hardware verified for Driver
- RDL+Bump Capability at Bump OSAT for 12-inch driver

SUBTASKS (7)

- Execute Mechanical wafers shortflow for FET (BEOL masks)
- Execute Mechanical wafers shortflow for Driver (BEOL masks, RDL , BUMP Masks)
- Verify Probe Touchdown on 12 inch wafers
- Embed FET mechanical wafers
- Characterize FA, Warpage and process
- Complete and sign-off SWIFT Gate 2 Deliverables for the Module
- Complete and sign-off SWIFT Gate 3 Deliverables for the Module

ASSUMPTIONS (3)

- Driver Node X Mech wafers , Target FET Fab Mech Wafers, Bump Mask and Stencil TO complete at Bump OSAT
- Capacity is secured with Mid-end OSAT and Bump OSAT and manufacturing flow is frozen
- 12inch FET development and Driver Node X driver development are sufficiently advanced to do BEOL tape-outs to generate Mechanical wafers

RISKS (5)

- Glass bonding process yield and reliability unknown - new for this product
- FSM/BGBM 12-inch schedules at Mid-end OSAT still TBD - late start could delay module builds
- Equipment qualification at new sites may uncover unexpected process gaps
- Multiple new process steps simultaneously increase integration risk
- if unable to secure capacity , then mid end and back end manufaturing flow could change

M2 Driver Node X Process Flow

Driver Node X Process Flow Frozen

DONENESS (4)

- Production ready PDK available
- Driver Node X foundry running HPC Lite technology qualification lots with high confidence to finish and pass qual by 1Q27
- Driver Node X HPC lite process flow fully ready with manufacturing controls, SPC , WAT coverage
- Driver Node X allowing fabrication and stock up of 12 inch full flow FET wafers for Project Beta Module Development

SUBTASKS (5)

- Release PDK
- Freeze process
- Tech Qual lot and qual execution
- Conduct Process readiness review with PG
- Tech production release safe launch Plan

ASSUMPTIONS (1)

- high confidence process freeze with no major changes to process or PDK expected

RISKS (2)

- Tech qual fails could drive major changes to process or PDK , this will impact Driver die B0 tapeout and Module schedule
- 5V CMOS unable to achieve PDK -0 maurity in time

M3 Gen 3 FET 12-inch Target FET Fab

Gen 3 SGMOS FET process frozen

DONENESS (3)

- Legacy FET Fab to Target FET Fab Process Transfer complete - parametric and wafer level reliability equiavlence proven
- All tooling gaps closed for Target FET Fab to autonomously support production volumes
- Technology Transfer Phase Gate process executed with evidence

SUBTASKS (7)

- Tool up Target FET Fab fab for Gen 3 SGMOS process leverage Reference Tool Fab tools for gaps
- Tapeout eval vehicle and run learning cycles (by Q1 2026)
- Tapeout SGMOS3 (Q4 202- and generate POR wafers
- WAT and Product yield meet yield and parametric match entitlement criterion
- Driver probe ready for yield and parametric data review
- Develop plan to Close remaining tooling gaps for autonomous production support
- Conduct Technology safe launch plan

ASSUMPTIONS (2)

- CAPEX approval and investment for Target FET Fab tooling
- Legacy FET Fab to Target FET Fab transfer tracking to committed plan from IM

RISKS (5)

- Significant differences in performance between Legacy FET Fab and Target FET Fab FETs requiring more process tuning learning cycles than anticipated
- Gen 3 FETs may not meet expected metrics on 12-inch, particularly with 50-micron backgrind
- CAPEX approval delay for tooling investment
- Tooling gaps between Target FET Fab and Reference Tool Fab may be larger than expected
- Any unanticipated delays in process transfer

MAJOR MILESTONES

Milestones M4–M6 · M-id, summary task, doneness criteria, subtasks, assumptions, and risks

Generic re-skin

Source: Project Beta Macro workbook

Major Milestones · Page 3 of 4

M4 Lookahead Module Build

Lookahead module with 8-inch FETs + 12-inch Driver Node X driver built, validated and look ahead qualified

DONENESS ⁽⁴⁾

- Functional Module build , Mid End and Back-end processing
- (8inch Legacy FET Fab + Reference Wafer Source is preferred path)
- Module Build Validation and Look ahead qualification
- Learnings/ Errata related to driver die funneled back to driver IC design team

SUBTASKS ⁽¹²⁾

- Source 8-inch FET wafers from Legacy FET Fab.and driver die from Driver Foundry
- Validate Driver die A0 silicon (WB package)
- Substrate readiness (stencil changes) to enable Driver Node X driver
- Embedded Laminate ready and staged at Assy
- Assemble lookahead modules using Driver Node X driver + 8-inch FETs
- Characterize and validate module
- RCM with right maturity level for this variant
- Create qual plan
- Qualify and validate hardware sourced and ready
- Complete SWIFT Gate 4 deliverables for Driver IC A0
- Complete SWIFT Gate 4 deliverables for Module for 8inch FET + Driver Node X A0 Driver
- Complete SWIFT Gate 2 and Gate 3 deliverables for Driver IC

ASSUMPTIONS ⁽⁵⁾

- 1PCB POD is FET wafer size agnostic - 8-inch/12-inch mix is viable
- Driver Node X HPC Lite 5V CMOS PDK reaches sufficient maturity (PDK V0.x?)
- Driver Node X driver silicon available from before Target FET Fab FETs
- Project Alpha team is able to spare some 8 inch laminates for Project Beta
- team to use 8inch FET + Driver Node X Driver is part of Project Beta scope. Both variants will be released simulatnaeously

RISKS ⁽¹⁾

- Driver Node X process is not fully matured (final PDK) by A0 tapeout resulting in dirver performance differences

M5 12-inch Module Flow

Target FET Fab 12 inch production flow with A0 driver built

DONENESS ⁽⁵⁾

- Functional Module build on full 12inch flow
- Process CZ reports from all OSATs on inline data
- Packaging updates PG on remaining risk /process improvements
- Module Validation and Lookahead qualification with A0 driver
- End-to-end cycle time and yield targets demonstrated

SUBTASKS ⁽⁴⁾

- Fab Driver Node X Driver die protolot and bump
- Fab FET protolot through Target FET Fab, FSM, BGBM and embedding
- Conduct SWIFT Gate 3 review for Driver B0 Die
- Conduct Module Gate 4 review for 12 inch FET + Driver Node X A0 Driver

ASSUMPTIONS ⁽¹⁾

- Any process optimizations for Mid End and Back-end from M2 are in place before running M5

RISKS ⁽⁴⁾

- Glass bonding process yield and reliability unknown - new for this product
- FSM/BGBM 12-inch schedules at Mid-end OSAT still TBD - late start could delay module builds
- Equipment qualification at new sites may uncover unexpected process gaps
- Multiple new process steps simultaneously increase integration risk

M6 Project Beta Module Build

Modules with 12-inch FETs + B0 Driver Node X driver built

DONENESS ⁽⁵⁾

- B0 G55 drive tape out completed
- 12-inch Target FET Fab FETs completed through post-fab processing (FSM/probe/BGBM) in final mfg flow
- Embedded substrate processed through final mfg flow frozen
- B0 Driver Node X driver probed/BGGM through final mfg flow
- Module assembled through final mgf flow (POR flow frozen)

SUBTASKS ⁽⁷⁾

- Tape out FETs to Target FET Fab-only process
- Determine scope of B0 changes based on A0 results
- Tape out B0 driver to Driver Foundry
- Qualify post-fab flow and assembly for both Target FET Fab and Driver Node X
- Complete post-fab processsing for both Target FET Fab FETs/Legacy FET Fab and Bo Driver Node X driver
- Assemble modules through qualified mfg flow
- Complete SWIFT Gate 2 and Gate 3 deliverables for B0 Driver IC

ASSUMPTIONS ⁽³⁾

- Target FET Fab and Driver Node X flows are frozen
- 12-inch module flow proven out
- Performance equivalence to Project Alpha is a pass/fail requirement

RISKS ⁽⁴⁾

- DC thermal performance may not match Project Alpha
- Mold void under IC driver (DOE and mold flow simulation required)
- Capacitor cracking or delamination during reflow
- First-time integration of all 12-inch components

M7 Customer Samples

Customer samples shipped (ES or above)

DONENESS (4)

- E1 validation completed
- LAR (or equivalent risk release) completed
- Review data and get sample shipment approval
- Engineering samples delivered to Customer Tier (Quantity TBD)

SUBTASKS (2)

- Determine sample quantity and plan builds accordingly
- Complete SWIFT module Gate 4 deliverables

ASSUMPTIONS (1)

- Project Alpha and Project Beta show no differences per internal validation

RISKS (1)

- Might need to increase risk builds if yields are not up to standards at this point

M8 Qual & Validation

Legacy FET Fab and Target FET Fab FETs qualified and validated

DONENESS (4)

- Full qual completed for both POR flows per agreed upon qual plan
- Si bench B0 validation completed across full operating conditions, temperature, and voltage range per RCM
- ATE test program 100% complete and characterization completed across temperature and voltage per RCM
- RCM fully satisfied and equivalent to Project Alpha

SUBTASKS (8)

- Develop Qual test program with transient completed
- Seup ENG Handler
- Conduct stress tests, review test data and analyze delta
- Develop characterization test program
- Characterize and review data
- Validate (full bench)
- Conduct Characterization and Validation reviews
- Complete SWIFT module Gate 4 and 5 deliverables

ASSUMPTIONS (4)

- End Customer and AI task force qual requirements are included in qual plan
- Manufacturing flow differences (12-inch wafers, glass bonding vs Taiko, different fab sites) drive incremental quals
- Qual hardware is available from IQT/LAR activities
- ELFR study for 20 dppm will not gate release - follows after MP release

RISKS (3)

- Qualification failure requiring design or process changes would add significant delay
- Reliability failure on embedded substrate
- IMON accuracy may not meet PRD specs (new trimming procedures under investigation)

M9 PCN Issued & Product Release

PCN issued to customers and Project Beta released to market

DONENESS (5)

- Process Change Notification submitted to Customer Tier
- Manufacturing flow, test programs, and supply chain locked
- Product declared mass production ready
- Datasheets published and part numbers orderable
- Production stocking in place

SUBTASKS (5)

- Prepare PCN documentation covering Project Alpha to Project Beta differences
- Submit PCN to Customer Tier for approval
- Lock test programs
- Complete SWIFT module Gate 6 deliverables
- Publish datasheets and enable ordering

ASSUMPTIONS (2)

- Overlap period planned: both Project Alpha and Project Beta in production before Project Alpha phase out
- Safe launch and ELFR study will not gate release - post-MP activities

RISKS (2)

- Customer PCN approval delay or rejection
- Customer requires additional testing or data beyond standard PCN package

MAJOR MILESTONES

Project Beta · M1–M9 · M-id, summary task, doneness criteria, subtasks, assumptions, and risks · generic re-skin

Source: Project Beta Macro workbook · 9 milestones

M#	Summary Task	Title	Doneness Criteria	Subtasks	Assumptions	Risks
M1	12-inch Module Flow	Mid-End Capability for 12 inch (post fab processing FET-FSM, BGBM , Driver - RDL/Bump)	<ul style="list-style-type: none"> 8/12-inch Target FET Fab FSM capability established at Mid-end OSAT 8/12-inch Target FET Fab BG/BM with glass bonding proven out at Mid-end OSAT 12-inch probe hardware verified for FET 12 -inch probe hardware verified for Driver RDL+Bump Capability at Bump OSAT for 12-inch driver 	<ul style="list-style-type: none"> Execute Mechanical wafers shortflow for FET (BEOL masks) Execute Mechanical wafers shortflow for Driver (BEOL masks, RDL , BUMP Masks) Verify Probe Touchdown on 12 inch wafers Embed FET mechanical wafers Characterize FA, Warpage and process Complete and sign-off SWIFT Gate 2 Deliverables for the Module Complete and sign-off SWIFT Gate 3 Deliverables for the Module 	<ul style="list-style-type: none"> Driver Node X Mech wafers , Target FET Fab Mech Wafers, Bump Mask and Stencil TO complete at Bump OSAT Capacity is secured with Mid-end OSAT and Bump OSAT and manufacturing flow is frozen 12inch FET development and Driver Node X driver development are sufficiently advanced to do BEOL tape-outs to generate Mechanical wafers 	<ul style="list-style-type: none"> Glass bonding process yield and reliability unknown - new for this product FSM/BGBM 12-inch schedules at Mid-end OSAT still TBD - late start could delay module builds Equipment qualification at new sites may uncover unexpected process gaps Multiple new process steps simultaneously increase integration risk if unable to secure capacity , then mid end and back end manufacturing flow could change
M2	Driver Node X Process Flow	Driver Node X Process Flow Frozen	<ul style="list-style-type: none"> Production ready PDK available Driver Node X foundry running HPC Lite technology qualification lots with high confidence to finish and pass qual by 1Q27 Driver Node X HPC lite process flow fully ready with manufacturing controls, SPC , WAT coverage Driver Node X allowing fabrication and stock up of 12 inch full flow FET wafers for Project Beta Module Development 	<ul style="list-style-type: none"> Release PDK Freeze process Tech Qual lot and qual execution Conduct Process readiness review with PG Tech production release safe launch Plan 	<ul style="list-style-type: none"> high confidence process freeze with no major changes to process or PDK expected 	<ul style="list-style-type: none"> Tech qual fails could drive major changes to process or PDK , this will impact Driver die B0 tapeout and Module schedule 5V CMOS unable to achieve PDK -0 maurity in time
M3	Gen 3 FET 12-inch Target FET Fab	Gen 3 SGMOS FET process frozen	<ul style="list-style-type: none"> Legacy FET Fab to Target FET Fab Process Transfer complete - parametric and wafer level reliability equiavlence proven All tooling gaps closed for Target FET Fab to autonomously support production volumes Technology Transfer Phase Gate process executed with evidence 	<ul style="list-style-type: none"> Tool up Target FET Fab fab for Gen 3 SGMOS process leverage Reference Tool Fab tools for gaps Tapeout eval vehicle and run learning cycles (by Q1 2026) Tapeout SGMOS3 (Q4 202- and generate POR wafers WAT and Product yield meet yield and parametric match entitlement criterion Driver probe ready for yield and parametric data review Develop plan to Close remaining tooling gaps for autonomous production support Conduct Technology safe launch plan 	<ul style="list-style-type: none"> CAPEX approval and investment for Target FET Fab tooling Legacy FET Fab to Target FET Fab transfer tracking to committed plan from IM 	<ul style="list-style-type: none"> Significant differences in performance between Legacy FET Fab and Target FET Fab FETs requiring more process tuning learning cycles than anticipated Gen 3 FETs may not meet expected metrics on 12-inch, particularly with 50-micron backgrind CAPEX approval delay for tooling investment Tooling gaps between Target FET Fab and Reference Tool Fab may be larger than expected Any unanticipated delays in process transfer
M4	Lookahead Module Build	Lookahead module with 8-inch FETs + 12-inch Driver Node X driver built, validated and look ahead qualified	<ul style="list-style-type: none"> Functional Module build , Mid End and Back-end processing (8inch Legacy FET Fab + Reference Wafer Source is preferred path) Module Build Validation and Look ahead qualification Learnings/ Errata related to driver die funneled back to driver IC design team 	<ul style="list-style-type: none"> Source 8-inch FET wafers from Legacy FET Fab.and driver die from Driver Foundry Validate Driver die A0 silicon (WB package) Substrate readiness (stencil changes) to enable Driver Node X driver Embedded Laminate ready and staged at Assy Assemble lookahead modules using Driver Node X driver + 8-inch FETs Characterize and validate module RCM with right maturity level for this variant Create qual plan Quality and validate hardware sourced and ready Complete SWIFT Gate 4 deliverables for Driver IC A0 Complete SWIFT Gate 4 deliverables for Module for 8inch FET + Driver Node X A0 Driver Complete SWIFT Gate 2 and Gate 3 deliverables for Driver IC 	<ul style="list-style-type: none"> 1PCB POD is FET wafer size agnostic - 8-inch/12-inch mix is viable Driver Node X HPC Lite 5V CMOS PDK reaches sufficient maturity (PDK V0.x?) Driver Node X driver silicon available from before Target FET Fab FETs Project Alpha team is able to spare some 8 inch laminates for Project Beta team to use 8inch FET + Driver Node X Driver is part of Project Beta scope. Both variants will be released simulatnaeously 	<ul style="list-style-type: none"> Driver Node X process is not fully matured (final PDK) by A0 tapeout resulting in driver performance differences
M5	12-inch Module Flow	Target FET Fab 12 inch production flow with A0 driver built	<ul style="list-style-type: none"> Functional Module build on full 12inch flow Process CZ reports from all OSATs on inline data Packaging updates PG on remaining risk /process improvements Module Validation and Lookahead qualification with A0 driver End-to-end cycle time and yield targets demonstrated 	<ul style="list-style-type: none"> Fab Driver Node X Driver die protolot and bump Fab FET protolot through Target FET Fab, FSM, BGBM and embedding Conduct SWIFT Gate 3 review for Driver B0 Die Conduct Module Gate 4 review for 12 inch FET + Driver Node X A0 Driver 	<ul style="list-style-type: none"> Any process optimizations for Mid End and Back-end from M2 are in place before running M5 	<ul style="list-style-type: none"> Glass bonding process yield and reliability unknown - new for this product FSM/BGBM 12-inch schedules at Mid-end OSAT still TBD - late start could delay module builds Equipment qualification at new sites may uncover unexpected process gaps Multiple new process steps simultaneously increase integration risk
M6	Project Beta Module Build	Modules with 12-inch FETs + B0 Driver Node X driver built	<ul style="list-style-type: none"> B0 G55 drive tape out completed 12-inch Target FET Fab FETs completed through post-fab processing (FSM/probe/BGBM) in final mfg flow Embedded substrate processed through final mfg flow frozen B0 Driver Node X driver probed/BGGM through final mfg flow Module assembled through final mgf flow (POR flow frozen) 	<ul style="list-style-type: none"> Tape out FETs to Target FET Fab-only process Determine scope of B0 changes based on A0 results Tape out B0 driver to Driver Foundry Qualify post-fab flow and assembly for both Target FET Fab and Driver Node X Complete post-fab processing for both Target FET Fab FETs/Legacy FET Fab and Bo Driver Node X driver Assemble modules through qualified mfg flow Complete SWIFT Gate 2 and Gate 3 deliverables for B0 Driver IC 	<ul style="list-style-type: none"> Target FET Fab and Driver Node X flows are frozen 12-inch module flow proven out Performance equivalence to Project Alpha is a pass/fail requirement 	<ul style="list-style-type: none"> DC thermal performance may not match Project Alpha Mold void under IC driver (DOE and mold flow simulation required) Capacitor cracking or delamination during reflow First-time integration of all 12-inch components
M7	Customer Samples	Customer samples shipped (ES or above)	<ul style="list-style-type: none"> E1 validation completed LAR (or equivalent risk release) completed Review data and get sample shipment approval Engineering samples delivered to Customer Tier (Quantity TBD) 	<ul style="list-style-type: none"> Determine sample quantity and plan builds accordingly Complete SWIFT module Gate 4 deliverables 	<ul style="list-style-type: none"> Project Alpha and Project Beta show no differences per internal validation 	<ul style="list-style-type: none"> Might need to increase risk builds if yields are not up to standards at this point
M8	Qual & Validation	Legacy FET Fab and Target FET Fab FETs qualified and validated	<ul style="list-style-type: none"> Full qual completed for both POR flows per agreed upon qual plan Si bench B0 validation completed across full operating conditions, temperature, and voltage range per RCM ATE test program 100% complete and characterization completed across temperature and voltage per RCM RCM fully satisfied and equivalent to Project Alpha 	<ul style="list-style-type: none"> Develop Qual test program with transient completed Seup ENG Handler Conduct stress tests, review test data and analyze delta Develop characterization test program Characterize and review data Validate (full bench) Conduct Characterization and Validation reviews Complete SWIFT module Gate 4 and 5 deliverables 	<ul style="list-style-type: none"> End Customer and AI task force qual requirements are included in qual plan Manufacturing flow differences (12-inch wafers, glass bonding vs Taiko, different fab sites) drive incremental quals Qual hardware is available from IQT/LAR activities ELFR study for 20 dppm will not gate release - follows after MP release 	<ul style="list-style-type: none"> Qualification failure requiring design or process changes would add significant delay Reliability failure on embedded substrate IMON accuracy may not meet PRD specs (new trimming procedures under investigation)
M9	PCN Issued & Product Release	PCN issued to customers and Project Beta released to market	<ul style="list-style-type: none"> Process Change Notification submitted to Customer Tier Manufacturing flow, test programs, and supply chain locked Product declared mass production ready Datasheets published and part numbers orderable Production stocking in place 	<ul style="list-style-type: none"> Prepare PCN documentation covering Project Alpha to Project Beta differences Submit PCN to Customer Tier for approval Lock test programs Complete SWIFT module Gate 6 deliverables Publish datasheets and enable ordering 	<ul style="list-style-type: none"> Overlap period planned: both Project Alpha and Project Beta in production before Project Alpha phase out Safe launch and ELFR study will not gate release - post-MP activities 	<ul style="list-style-type: none"> Customer PCN approval delay or rejection Customer requires additional testing or data beyond standard PCN package

PROGRAM MILESTONE TIMELINE

Fab1 | Foundry Partner 1B DRAM + DRAM+ (FeCAP) | Region B | 14,000 sqm GMA | Project start: Jan 1, 2027



KEY METRICS

Exec Sponsor: Exec Sponsor / CEO | CEO: Exec Sponsor / CEO | Procurement: Procurement Lead | Finance: Finance Lead | EPC: EPC Partner (Reference Fab A / Reference Fab B reference)
 Foundry: Foundry Partner 1B DRAM license | Process: Classical DRAM + DRAM+ FeCAP | Capacity: 20,000 WSPM Phase 1, 100,000 WSPM Phase 2 | Funding: \$4.9B (Sovereign Chips Grant \$2.875B / 58%)
 Pre-project status (April 2026): site undecided (FAB-01), Foundry Partner unsigned (FAB-02), EPC Partner design not started (FAB-03) | Project start M0 = Jan 1, 2027 (Site Selected & Acquired) | First Revenue M44 = Sep 2030

MILESTONE DETAILS

Fab1 | 13 milestones | Site Selected & Acquired through DRAM+ (FeCAP) qualified

M1 Jan '27

Site Selected & Acquired

Region A vs Region B resolved. Land acquisition closed. SPV established. All downstream gates now unblocked.

50 staff | Project kickoff | M0

M6 Jun '27

Groundbreaking

EPC Partner mobilisation. Earth works begin under partial permit (Partial Permit). Site cleared and foundations started.

200 staff | Construction mobilisation | M5

M11 May '30

Baseline DRAM qualified

Foundry Partner-licensed baseline DRAM process qualified on Fab1 tools. Gate to all downstream product runs.

400 staff | Process qual team | M40

M2 Feb '27

Foundry Partner license signed

Foundry Partner Technology license executed. Detailed tool list and process consumption data released to EPC Partner and Client Co.

30 staff | Foundry Partner transfer | M1

M7 May '28

Fab weathertight

Building shell closed. Off-Site Manufacturing (OSM) MEP modules begin install. Cleanroom fit-out commences.

600 staff | Peak civil + MEP | M16

M12 Sep '30

First revenue / Phase 1 qual

Phase 1 product qualification complete. First revenue from 20k WSPM Classical DRAM. \$300M to 900M per year by 2031.

300 staff | Production ramp | M44

M3 Mar '27

Environmental Permit issued

Full federal environmental permit (Environmental Permit) granted. Partial Permit enabled earth works at M2 ahead of full permit.

40 staff | Permitting + legal | M2

M8 Nov '28

Fab shell complete

Full structural and envelope completion. Cleanroom and sub-fabs ready for tool delivery.

800 staff | Pre-RFE peak | M22

M13 Mar '31

DRAM+ (FeCAP) qualified

Ferroelectric capacitor (FeCAP) DRAM+ technology integrated and qualified. \$1B+ per year by 2031 per Fab1 forecast.

350 staff | FeCAP integration | M50

M4 Apr '27

EPC Partner L0 design start

EPC Partner L0 Pre-Concept design contract issued and funded. Every month of L0 delay slips RFE one-for-one.

100 staff | EPC Partner design ramp | M3

M9 Dec '28

RFE — EPC Partner target

Ready For Equipment per EPC Partner 14,000 sqm GMA gantt. Achievable via OSM + Partial Permit. Stretch target.

1000 staff | Tool install ramp | M23

M5 May '27

Long-lead tool POs placed

Binding POs issued for Lithography Vendor lithography (28 tools) and ALD Vendor. 18 to 24 month lead times locked. \$800M committed.

50 staff | Procurement + finance | M4

M10 Jun '29

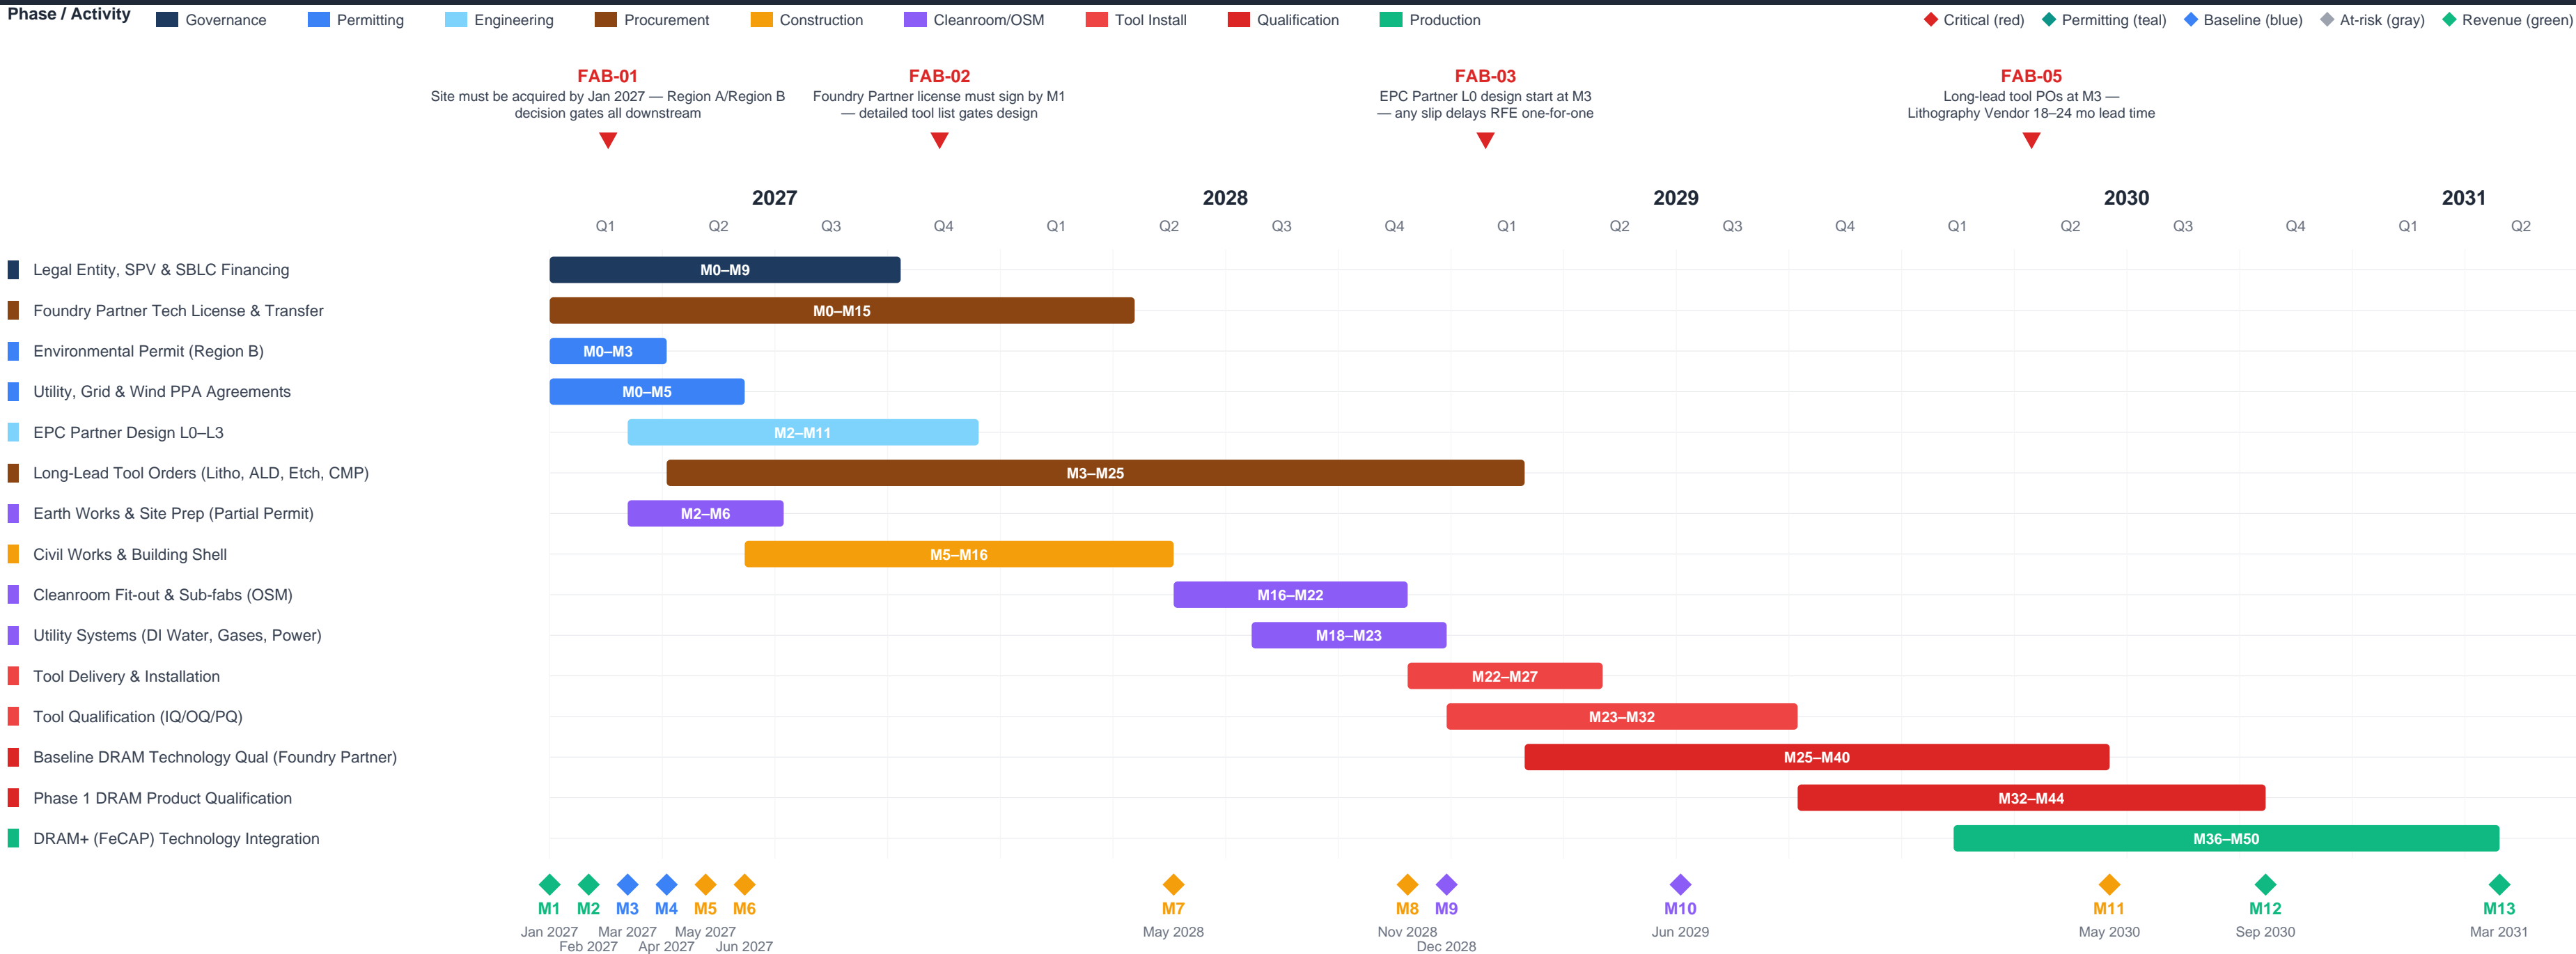
RFE — Federal Baseline

Federal Baseline RFE. 6 months later than EPC Partner target. Current optimistic scenario for qualification start.

1000 staff | Tool install peak | M29

FAB1 — MACRO SCHEDULE WITH RISK ANALYSIS

Fab1 | Foundry Partner 1B DRAM + DRAM+ (FeCAP) | Region B | 14,000 m² GMA | \$4.9B | Project start: Jan 1, 2027 (Site Selected & Acquired)



Pre-project status as of April 2026 (9 months before M0)

- **SITE:** Region A preferred but land not acquired (April 2026). City refuses technical discussions until financial resolution. Region B active since Nov 2025. Target acquisition: Jan 2027 (M0).
- **FOUNDRY PARTNER:** License not yet signed. Expected Apr/May 2026. Detailed tool list and consumption data only available post-signing. Must close before M1 (Feb 2027).
- **EPC PARTNER DESIGN:** Not started. Monthly run rate for site support only. Fab design order gated on site + Foundry Partner. Must start by M2 (Mar 2027) to preserve schedule.
- **IMPACT:** 8-month slip from previous Q1 2026 baseline. New project anchors at Jan 2027 = M0 (Site Selected & Acquired).
- **CRITICAL PATH:** Site M0 → Foundry Partner M1 → Environmental Permit M2 → EPC Partner L0 M3 → Groundbreaking M5 → Weathertight M16 → RFE M23–M29 → First

Financing Structure (\$4.9B Total)

- Sovereign Chips Grant: \$2.875B (~58%) — up from \$1.3B Jan 2025 indication.
- PE Sponsor PE equity via SPV; \$4.9B SBLC from bank consortium (Lead Bank lead).
- SBLC expires at First Silicon — drawdown structured against Sovereign Chips Grant milestones.
- OPEX buffer \$302M; SBLC interest provision \$350M.
- EPC Partner team: Reference Fab A + Reference Fab B reference; OSM reduces on-site schedule risk.

Schedule Scenarios & Capacity

- EPC Partner scenario: RFE M23 (Dec 2028), First Silicon M32 (Sep 2029). OSM + Partial Permit enabled.
- Federal Baseline: RFE M29 (Jun 2029), production start M32 (Sep 2029) → First revenue M44 (Sep 2030).
- Phase 1: 20,000 WSPM — 14,000 m² GMA cleanroom (EPC Partner design basis).
- Phase 2: 100,000 WSPM — 2nd-gen DRAM/DRAM+ ('Developed in Sovereign Bloc') from 2033.
- LOI customers: Anchor OEMs + 1,000+ Sovereign Bloc OEMs.

Qualification Sequence & Revenue Forecast

- Baseline Foundry Partner DRAM qual (M25–M40): gate to all downstream — runs from RFE.
- Phase 1 product qual (M32–M44): first revenue Sep 2030 (\$300M–900M/yr by 2031).
- DRAM+ (FeCAP) integration (M36–M50): \$1B+/yr by 2031 per Fab1 forecast.
- Phase 2 (100k WSPM): 2nd-gen DRAM/DRAM+ from 2033, \$3B+/yr by 2034.
- Baseline DRAM revenue de-risks SBLC and business case during FeCAP ramp.

PROGRAM CONTEXT

Fab1 | Success criteria, key assumptions, and top risks | Sourced from Fab1 Project Narrative

SUCCESS CRITERIA

- Site decision finalised (Region A or Region B confirmed) and land acquisition completed by M6.
- Sovereign Chips Grant (\$2.875B, ~58%) committed and disbursement schedule agreed with Federal Ministry.
- Foundry Partner technology license executed by M4 to M5 (Apr/May 2026); detailed tool list and consumption data received.
- EPC Partner fab design contract issued immediately after site + Foundry Partner confirmed; L0 no later than M6.
- \$4.9B SBLC from bank consortium closed; PE Sponsor PE equity committed.
- Groundbreaking (EPC Partner mobilisation) by M9 (Q4 2026); fab shell complete by M26.
- Partial Permit enabling earth works at M3; fab weathertight by Q3 2027.
- Ready For Equipment (RFE): M27 (Apr 2028) EPC Partner target, M33 (Q4 2028) Federal Baseline.
- Full IQ/OQ/PQ completed within 3 months of RFE.

KEY ASSUMPTIONS

- Site decision reached by M6: either Region A financial impasse resolved or Region B plot confirmed and acquired.
- Sovereign Chips Grant (\$2.875B) disbursement aligns with construction cash-flow requirements.
- Foundry Partner license signs by end April / May 2026 (M4 to M5); detailed tool list and consumption data available.
- EPC Partner fab design contract issued and funded by M6; L0 start no later than M6 to preserve Federal Baseline.
- Partial Permit obtainable by M3 to enable early earth works before full Environmental Permit.
- ALD Vendor tooling available at required volume and delivery schedule; EPC Partner manages tool install.

KEY RISKS

- CRITICAL**
Site undecided at M4: Region A blocked on financials; Region B alternative; no land acquired. Gates all permitting and design.
- CRITICAL**
Foundry Partner license not signed at M4; detailed tool list and consumption data unavailable; EPC Partner design cannot start without these inputs.
- CRITICAL**
EPC Partner design not started at M4; every month of delay to L0 start directly slips RFE. Federal Baseline M33 requires design start by M6.
- HIGH**
Lithography Vendor DUV lead times 18 to 24 months; POs cannot be placed without Foundry Partner tool list; already late vs M6 target.
- HIGH**
SBLC bank consortium does not close at required terms — construction cannot begin.
- HIGH**
ALD Vendor single-source dependency: slot unavailability delays process qual by 6+ months.
- HIGH**
EPC Partner: construction cost overrun or schedule slip beyond contract completion date.
- MEDIUM**
FeCAP novel technology: no production-scale precedent; yield ramp M44 to M54 uncertain.
- MEDIUM**
CAPEX inflation: current \$4.9B estimate is a 49% increase vs Jan 2025 \$3.3B baseline.
- MEDIUM**
Foundry Partner financial commitment not provided; Federal Ministry condition on Foundry Partner equity not resolved.
- LOW**
Wind park PPA not secured; utility agreement gates Environmental Permit application.
- DRAM market downturn 2026 to 2029 weakens LOI conversion to binding offtake contracts.

SCHEDULE NOTES — ANALYSIS & LOGIC

Supporting analysis, schedule logic, qualification sequence, and open gates as of April 2026 (M4)

Schedule Logic — Status as of April 2026 (M4)

As of April 17, 2026 (M4), three design gates remain open: site is undecided (Region A vs Region B), the Foundry Partner license is not signed, and EPC Partner fab design has not started. The EPC Partner September 2025 proposal (RFE April 2028, M27) assumed L0 design start in October 2025 — this never happened. The Federal Baseline (RFE Q4 2028, M33) is the current optimistic scenario, achievable only if Foundry Partner signs by M4–M5, site is confirmed by M6, and EPC Partner L0 starts no later than M6–M7. Each month of further delay slips RFE by one month.

Open Gates as of April 2026 (M4) — Site Lead, Client Co

SITE (Site Lead, April 17 2026): Region A selected but land not acquired. City blocked technical discussions until financial resolution; Client Co waiting. Region B contacted Client Co Nov 2025. Site 'not 100% closed.'

FOUNDRY PARTNER: License expected end Apr / early May 2026. Only post-signing does Client Co receive detailed tool list + consumption data — the two remaining design inputs.

EPC PARTNER DESIGN: Never started. Monthly run rate for site-selection support only. No fab design order issued. The EPC Partner L0→RFE schedule in the Sep 2025 workshop is a proposal, not active work.

EPC Partner Construction Approach (Sep 2025 Workshop)

EPC Partner's own project Gantt (14,000 m² GMA / 20k WSPM) targets RFE April 2028 (M27), 6 months faster than the Federal Baseline of Q4 2028 (M33). Key enablers:

- Partial Permit: Earth works start M3, before full Environmental Permit at M6. This is standard German staged permitting (Staged Permit) — not in prior assumptions.
- Off-Site Manufacturing (OSM): Pre-fabricated MEP modules assembled off-site in parallel with on-site construction, reducing peak on-site workforce and compressing schedule.

Critical Path Items

Tool procurement must begin at M4–M6 regardless of construction status. Lithography Vendor DUV lithography tools carry 18–24 month lead times. The Foundry Partner technology license must also be executed early — the DRAM process recipe transfer runs in parallel with construction and must be ready before RFE at M33.

The Environmental Permit target of M6 (Q3 2026) is significantly more aggressive than the 12–18 month German norm. The Region B site's existing permitting alignment (Federal Baseline already issued) is the key enabler. Any 3-month slip compresses the construction programme and threatens the Q1 2029 First Silicon target.

Qualification Sequence & Revenue Forecast

Per Fab1 revenue forecast (Federal Baseline slide 7):

- Baseline DRAM qual (M33–M44): Foundry Partner-licensed process on Fab1 tools. Gate to all downstream. Estimated \$300M/yr at full Phase 1 utilisation.
- Phase 1 DRAM product qual (M38–M50): revenue product qualification. First revenue ~2030 (\$300M–900M/yr by 2031).
- DRAM+ / FeCAP technology (M44–M54): novel technology, \$1B+/yr by 2031 per Fab1 forecast. Phase 2 (100k WSPM) adds \$3B+/yr from 2034.

Fab1-Specific Considerations

EPC Partner as of September 2025 stated: 'Target first wafer out timing is potentially at risk, mitigation measures are required.' The risk: Sovereign Chips Grant, PE, and Foundry Partner were all still pending at that point. The pre-FID design approach (L0–L2) is the primary mitigation — design work proceeds on benchmarks until Foundry Partner data is available.


Critical dependency: Foundry Partner must be on board before start of Concept (L2). If not, EPC Partner advises a 'project break could be considered to avoid costly rework.' EPC Partner scope extends beyond EPC basebuild to include tool install, qualifying the fab through to production-ready. EPC Partner team includes Reference Fab A and Reference Fab B experience — strongest Foundry Country-specific cleanroom construction reference available.

RISK REGISTER & CAPEX PROCUREMENT ANALYSIS

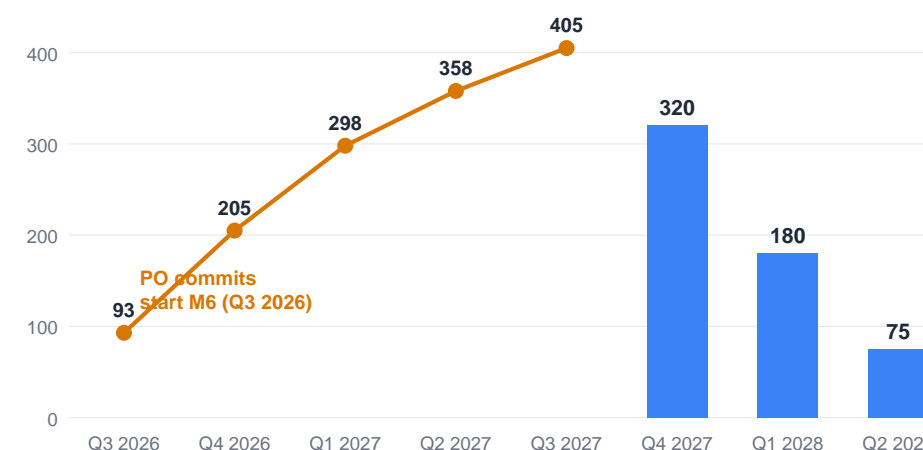
Top schedule & execution risks ranked by impact | Tools & Technology CAPEX \$2.41B | Procurement timeline

ID	Severity	Risk Description & Detail	Impact	Mitigation
FAB-01	CRITICAL	Site undecided at M4 (Region A blocked, Region B alternative) Region A: land not acquired; city blocks technical discussions until financial resolution. Region B under active evaluation since Nov 2025.	+6–18 mo	Force Region A financial resolution OR formally pivot to Region B. Decision required by M6 to preserve Federal Baseline schedule.
FAB-02	CRITICAL	Foundry Partner license not signed at M4; design inputs unavailable License expected end Apr / early May 2026. Detailed tool list + consumption data only released post-signing.	+4–12 mo	Close Foundry Partner license by end of April. Immediately extract detailed tool list and consumption data post-signing (target: 2 weeks).
FAB-03	CRITICAL	EPC Partner design not started at M4 — Federal Baseline M33 now at risk EPC Partner never contracted for fab design. Sep 2025 M27 target lost. Every month of further delay slips RFE.	+3–6 mo / month	Issue EPC Partner L0 design contract immediately after Foundry Partner signs + site confirmed. Target contract execution by M5–M6.
FAB-04	HIGH	Foundry Partner technology transfer incomplete at RFE DRAM process must be qualified on Fab1 tools by M33. Foundry Partner financial commitment not yet resolved by Federal Ministry.	+3–9 mo	Execute Foundry Partner license by M5. Station Client Co engineers at Foundry Partner Foundry Country for parallel process learning during construction.
FAB-05	HIGH	SBLC bank consortium does not close at required terms \$4.9B SBLC requires bank hold >50% LTC after Sovereign Chips Grant. Complex structure may delay beyond M9 target.	+6–18 mo	Engage Lead Bank and consortium banks immediately. Align LTC/Sovereign Chips Grant drawdown structure with SBLC terms. Weekly tracking.
FAB-06	HIGH	ALD Vendor single-source dependency at fab scale ALD Vendor (DRAM+): 1 tool. Already flagged as risk in DDR4 and Cache+ programs. Slot availability uncertain.	+3–9 mo	CEO-level ALD Vendor engagement. Validate fab-scale slot availability at Q2 2026. Backup evaluation for non-FeCAP.
FAB-07	HIGH	EPC Partner: construction or tool install overrun / delay EPC Partner responsible for basebuild + tool install to Fab-Ready. Delay compresses qualification timeline.	+3–6 mo	Tight contractual milestones with LDs. EPC Partner Gantt: M1 Kickoff to M14 Weathertight, M27 RFE.
FAB-08	HIGH	Foundry Partner not on board at Concept (L2) start — design break risk EPC Partner: if Foundry Partner not available at L2, 'project break could be considered to avoid costly rework.'	+2–4 mo	Foundry Partner license must be executed by M3. EPC Partner benchmarks fill gap if data arrives by L2 end (M5–M6).
FAB-08b	HIGH	Cleanroom construction labor scarcity in Foundry Country region Fewer cleanroom specialists than Asia. EPC Partner, Competing Fab, and battery plants compete for same labor.	+3–6 mo	EPC Partner locks specialist subcontractors at contract award. OSM approach reduces on-site peak demand.
FAB-09	MEDIUM	CAPEX inflation vs Jan 2025 baseline (+49%) Total has risen from \$3.3B to \$4.9B. Sovereign Chips Grant eligible costs up from \$2.2B to \$4.9B.	+0 mo (cost risk)	\$4.9B estimate includes contingency. Monitor EPC Partner EPC fixed-price structure. Lock procurement at M6.
FAB-10	MEDIUM	Foundry Partner financial commitment not provided (Federal Ministry condition) Federal Ministry requested Foundry Partner equity/commitment as condition of Sovereign Chips Grant profit/loss obligation.	+6–12 mo	Client Co developing Foundry Partner commitment proposal. Federal Baseline support requested at federal level.
FAB-11	MEDIUM	FeCAP novel technology extends yield ramp No production-scale FeCAP precedent. Yield ramp M44–M54 is already tight in financial plan.	+6–12 mo	Run FeCAP short-loop validation at Foundry Partner concurrently. Model yield vs. spec in financial plan.
FAB-12	MEDIUM	DRAM market downturn weakens LOI-to-offtake conversion LOI customers confirmed (Anchor OEMs + 1,000+). Binding contracts pending.	+0 mo (revenue risk)	Convert LOIs to binding offtake by M18. Prioritise Key Accounts (Anchor OEMs).

Tools & Technology CAPEX — Total: \$2.41B (Project total: \$4.9B)



Tool Procurement Timeline (\$M)



Risk Summary

- 3 CRITICAL**
Site (01), Foundry Partner (02), EPC Partner design (03)
- 6 HIGH**
Foundry Partner transfer, SBLC, ALD Vendor, EPC Partner EPC, L2 break, labor
- 4 MEDIUM**
CAPEX inflation, Foundry Partner equity, FeCAP, LOI conversion

Risk Status Overview — M4 (April 2026)

- > 3 critical gates open: site decision (Region A vs Region B), Foundry Partner license, EPC Partner fab design contract
- > EPC Partner M27 (Apr 2028) RFE target LOST — every month of further delay to L0 slips RFE one-for-one
- > Federal Baseline M33 (Q4 2028) baseline achievable only if Foundry Partner signs by M5 and EPC Partner L0 starts by M6–M7
- > Tool POs at M6 (Q3 2026) are critical: Lithography Vendor DUV 18–24 mo lead time; \$2.41B tools + \$223M Foundry Partner license

TOOL INVENTORY & PRIORITY MITIGATION ACTIONS

Complete tool inventory | Procurement timing | Mitigation actions ranked by schedule recovery potential

Tool	Area	Qty	PO	Delivery	\$ M	Lead Risk
ArFi Scanner + Track	Litho	14	M6	M18–M27	627	CRITICAL
ArF Dry Scanner + Track	Litho	2	M6	M18	42	HIGH
KrF Scanner + Track	Litho	12	M6	M18–M30	131	HIGH
ALD (standard)	Thin Film	5	M6–M8	M18–M27	60	HIGH
ALD Vendor (DRAM+)	Thin Film	1	M8	M21	12	CRITICAL
CVD (EPI, WCVD, multi)	Thin Film	39	M8–M12	M21–M30	225	MEDIUM
PVD	Thin Film	8	M8	M18–M30	44	MEDIUM
Dry Etch (Cond/Dielectric)	Etch	37	M8–M12	M18–M27	283	MEDIUM
Dry Etch (Ash)	Etch	7	M8–M12	M18–M27	11	LOW
CD-SEM	Metro	8	M8	M18–M27	79	MEDIUM
Implant	Impl	11	M8–M12	M18–M27	84	MEDIUM
Thermal Furnace	Impl	12	M10–M12	M18–M30	39	LOW
Thermal RTA	Impl	8	M8–M12	M18–M30	26	LOW
Wet Processing	Impl	14	M8–M12	M18–M30	27	LOW
Metrology (multi)	Impl	45	M10–M12	M21–M27	214	MEDIUM
Inspection (multi)	Thin Film	45	M8–M12	M18–M27	123	MEDIUM
CMP	CMP	14	M8–M12	M18–M27	61	MEDIUM
Sort / Functional Test	Test	35	M12	M21–M30	38	LOW
Plating	Thin Film	2	M12	M21–M30	11	LOW
Spin-On	Thin Film	6	M12	M21–M27	16	LOW
MES + Yield Mgmt SW	IT	2	M10	M18–M21	4	LOW
AMHS	Facilities	1	M12	M21–M24	2	LOW
TOTAL		~328 tools			\$2,159M	

Total Project CAPEX Breakdown

> Building, Site, Shell, Cleanroom, Utilities (EPC Partner EPC)	\$1783M
> Foundry Partner Technology License & Transfer	\$223M
> Advanced Technology Tools (DRAM+)	\$213M
> R&D CAPEX (Models + Tests)	\$119M
> OPEX buffer (construction phase)	\$302M
> SBLC interest provision	\$350M

TOTAL PROJECT CAPEX ~\$4.9B

(up) from \$3.3B in January 2025 baseline (+49%)

Priority Mitigation Actions — Ranked by schedule recovery potential

MIT-01	<p>Force site decision NOW: resolve Region A or commit to Region B Exec Sponsor / CEO / Procurement Lead · Target: M4–M5 (immediate) Region A: issue financial ultimatum or pivot. Region B: confirm plot, commission land survey. No design can start without site.</p>	Saves up to 18 mo
MIT-02	<p>Close Foundry Partner license by end of April 2026 (M4) Exec Sponsor / CEO · Target: April 30 / May 2026 (M4–M5) License expected end April. Immediately extract detailed tool list + consumption data post-signing. Issue EPC Partner contract within 30 days.</p>	Saves up to 12 mo
MIT-03	<p>Place Lithography Vendor lithography POs at M6 Procurement Lead / Procurement · Target: M6 — PO binding Place binding POs for all 28 litho tools. Negotiate staged delivery. Escalate via Sovereign Chips Grant.</p>	Saves up to 12 mo
MIT-04	<p>Close Sovereign Chips Grant application — Federal Ministry / Federal Baseline priority Finance Lead / Exec Sponsor / CEO · Target: M3–M6 — Sovereign Chips Grant \$2.875B Prerequisite for SBLC. Advance Stage 2 in parallel. Bridge for tool deposits.</p>	Saves up to 24 mo
MIT-05	<p>Close SBLC with bank consortium (Lead Bank lead) Finance Lead / Procurement Lead · Target: M4–M9 Engage Lead Bank and consortium. Align LTC/Sovereign Chips Grant drawdown structure with SBLC terms.</p>	Saves up to 18 mo
MIT-06	<p>Issue EPC Partner L0 design contract immediately after site + Foundry Partner confirm Procurement Lead / Finance Lead · Target: M5–M6 (latest) EPC Partner design NEVER STARTED. Issue L0 contract as soon as Foundry Partner signs and site confirmed. Every week of delay = one week of RFE slip.</p>	Saves up to 9 mo
MIT-07	<p>Secure ALD Vendor commitment at fab scale CTO / Exec Sponsor / CEO · Target: M4 — CEO-level engagement Validate fab-scale slot availability. Backup evaluation for non-FeCAP.</p>	Saves up to 9 mo
MIT-08	<p>Stagger specialty tool POs across 2027 Procurement / Procurement Lead · Target: M6–M12 — spread Q1–Q4 2027 Spread etch, CVD, implant POs across Q1–Q4 2027. Secure vendor commitments with deposits.</p>	Saves up to 6 mo
MIT-09	<p>Lock wind park PPA and utility agreements Procurement Lead / Legal · Target: M4–M6 — PPA concept exists Execute term sheet now. Utility agreements gate Environmental Permit application.</p>	Saves up to 12 mo
MIT-10	<p>Begin Foundry Partner DRAM short-loop validation at Foundry Partner facility Tech Transfer Lead / Tech Transfer Lead · Target: M6–M24 Validate Fab1 process parameters at Foundry Partner before fab ready. Model yield vs. spec in financial plan.</p>	Saves up to 12 mo