Methodology brief

How lateralworks uses AI

How the FTTM methodology applies AI to accelerate new product development

lateralworks · FTTM advisory

A plain-language account of where lateralworks applies AI in the planning and execution of new product development, how client data is handled, and where human judgment stays in control.

Prepared by	Date	For
lateralworks FTTM advisory	June 2026 Confidential	Companion to the council briefing

Table of contents

How lateralworks uses AI

Core thesis. lateralworks uses AI for one job: synthesis. It pulls many data streams into one organized structure that feeds the plans and decision models a human core team then scrubs, decides on, and owns. The model accelerates the work. It does not make the call.

Overview

Executive summary

In one word, lateralworks uses AI for synthesis. The methodology brings together many data streams: schedules and milestone data, status updates and meeting notes, decision inputs, assumptions, and economic parameters. It organizes them into one integrated structure. That structure becomes the input to the lateralworks tools, the plans, the decision models, and the cost-of-delay models the program runs on [1].

This matters because the bottleneck in new product development is rarely a shortage of data. It is the time it takes to consolidate scattered data into something a team can act on. AI compresses that step from days to minutes, and the consolidated picture is higher in quality than hand assembly produces. The field evidence is consistent on both points. Controlled studies of AI-assisted knowledge work report large reductions in time-to-draft and measurable quality gains [2, 4]. For a program committed to nine products in two years, that compression is a force multiplier on the schedule.

AI is applied at five points in the FTTM workflow: planning and tracking, decision-making, economic modeling, schedule evaluation, and status reporting. At every one of them the AI output feeds a human decision rather than making it. The core team reviews, challenges, and owns the result [1].

lateralworks matches the control to the sensitivity of the data. The default configuration does not use conversations to train models, and retention is short [6]. For more sensitive work, stronger postures are available, up to an isolated model that runs entirely inside a client firewall. This brief sets out each of these in plain language, so the program team knows what the tools do, what they do not do, and where the human stays in command.

01

The principle

Synthesis, not autonomy

The principle

Synthesis, not autonomy

The one thing the tools do

Strip away the labels and lateralworks uses AI for a single function. It reads many separate sources and returns one organized structure. A program generates data in fragments: a schedule in one place, a status note in another, a set of assumptions in someone's head, an economic parameter in a spreadsheet. On their own, fragments do not drive decisions. Assembled, they do. AI does the assembly.

Why synthesis is the bottleneck worth attacking

In most programs the slow step is not deciding. It is getting to the point where a decision can be made. Pulling the inputs together, reconciling them, and laying them out in a form the core team can scrub absorbs hours that should go to judgment. The same pattern shows up in the case-study record: synthesis that used to take a planner days now takes minutes, and the team spends the recovered time on the decision rather than on the paperwork that precedes it [1].

The structure of the work

The pattern is always the same. Many data streams flow into one synthesis pass. The pass produces an integrated structure. That structure feeds the lateralworks tools, the plans, the decision models, and the cost-of-delay models. The output then goes to a human core team that reviews it, decides, and owns the result. The figure below shows the full path.

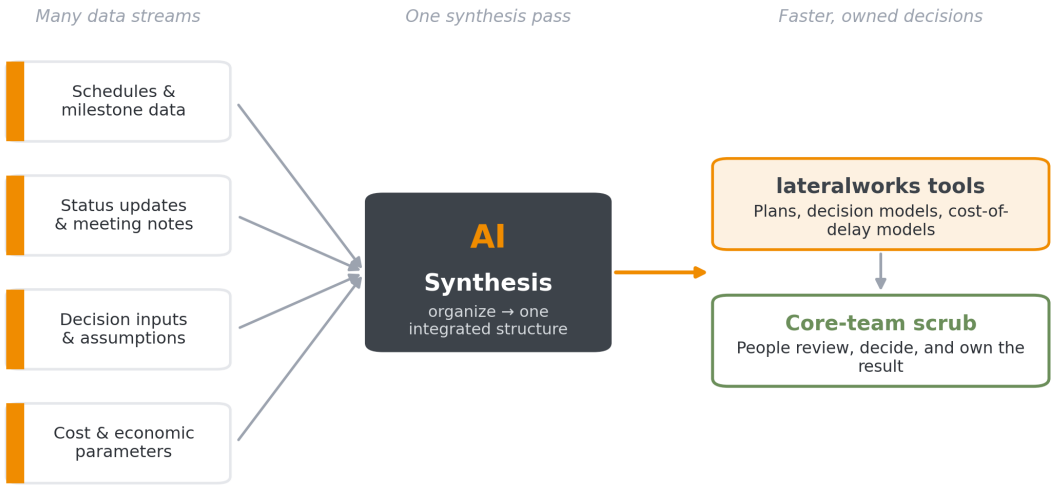

Figure 1. The synthesis pattern. AI consolidates many data streams into one structure that feeds the lateralworks tools. A human core team reviews, decides, and owns the outcome.

02

Where it works

Five points in the FTTM workflow

Where it works

Five points in the FTTM workflow

Planning and tracking

AI synthesizes the milestone timeline and the milestone matrix from raw inputs. Instead of hand-building the template from scattered schedule data, the team receives a consolidated draft and spends its time scrubbing it with the core members. This is the work already underway on the Fieldpiece program [1].

Decision-making

The methodology builds decision models and then interprets their results. AI helps construct the model from the inputs and helps read what the model is saying, so the team moves from question to a well-framed choice faster.

Economic modeling

Cost-of-delay is central to FTTM, and the economic models behind it are built and analyzed with AI support. A clear cost-of-delay picture is what lets a program prioritize correctly when everything feels urgent. It is the difference between nine products landing on time and a queue that slips.

Schedule evaluation

AI evaluates schedules to surface gaps and risks earlier than a manual read would. Problems that would otherwise appear at a milestone show up while there is still room to act on them.

Status reporting

Weekly project status reports are drafted from program data, returning hours to managers that would otherwise go to assembling updates by hand. The manager edits and owns the report. The draft is a faster starting point.

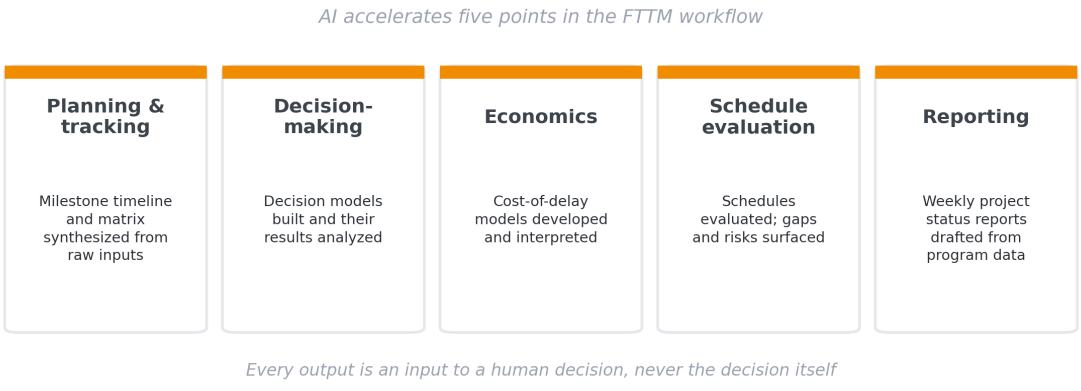

Figure 2. The five points where AI accelerates the FTTM workflow. At each one the AI output is an input to a human decision, never the decision itself.

03

Your data

Match the control to the data

Your data

Match the control to the data

The default lateralworks runs on

By default, lateralworks uses a configuration where the model provider does not use conversations to train its models. The training setting is off, and retention is short [6]. This is the posture already reviewed with Fieldpiece's contract team. For most program data it is the right balance of speed and protection.

Stronger postures, when the data calls for them

Data sensitivity is not one level, so the control should not be either. As data becomes more sensitive, the posture escalates. The consumer default gives way to commercial and enterprise terms, where training is contractually prohibited and retention tightens. Above that sits an API path with a data-processing agreement and zero or near-zero retention. At the top, an isolated model runs entirely inside the client firewall, so data never leaves. The figure shows the full ladder.

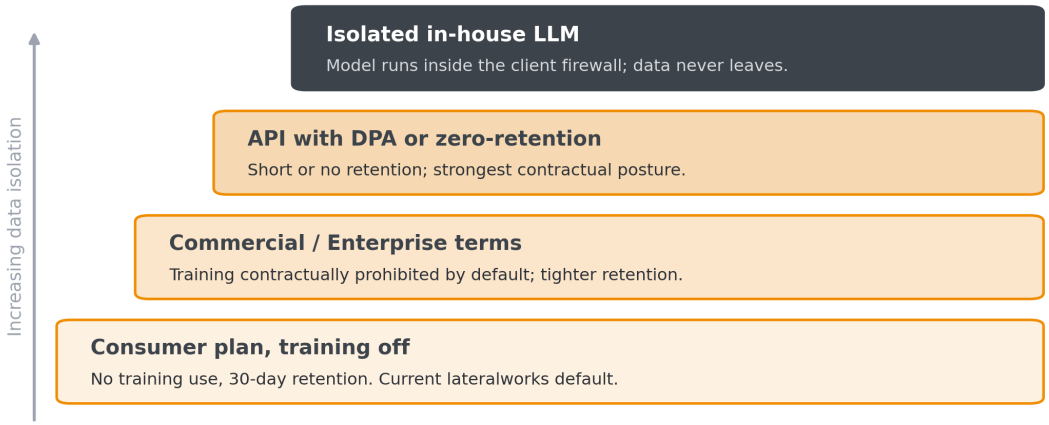

Figure 3. The data-isolation ladder. The control is matched to the sensitivity of the data, from the current lateralworks default up to a fully isolated in-house model.

The in-house path is real, not theoretical

The top rung is in active use. Through a strategic partner doing advanced AI work, lateralworks has helped a client stand up an internal LLM that sits inside their own firewall, applied to their development process to accelerate learning and shorten time to first product. When a client's data sensitivity warrants it, that capability is available rather than aspirational.

04

The boundary
What always stays human

The boundary

What always stays human

AI proposes; the core team disposes

Every application in this brief produces an input to a decision. None of them produces the decision. The synthesized plan, the model result, the cost-of-delay read, the schedule risk, the status draft: each lands in front of the core team, which scrubs it, challenges it, and signs off. Accountability stays with the people, where it belongs. The field evidence supports the design. The largest productivity studies find that AI helps most when the human keeps a clear role in reviewing and directing the work, and that errors rise when the boundary is lost [2, 3].

What the tools do not do

- They do not approve plans, commit dates, or accept risks. People do.
- They do not replace the core-team scrub. They prepare the material for it.
- They do not act on data outside the agreed handling posture for that data.
- They do not learn from client conversations under the default configuration [6].

Why this boundary is the point, not a caveat

Keeping the human in command is what makes the speed safe. A team that owns every decision can move fast on the synthesis because it knows the judgment still runs through it. Remove the human and the program has not gone faster. It has lost the accountability that lets it trust the pace. lateralworks is built the other way round: accelerate the assembly, protect the judgment.

> **The one-line summary.** lateralworks uses AI to consolidate many data streams into one structure that feeds its planning and decision tools, handled under a data posture matched to the data, and always scrubbed and owned by the human core team.

Sources

References

[1] lateralworks. “AI-assisted macro planning.” FTTM methodology series, Rev 1.1, 2026. Case-study record of the synthesis flow and the core-team scrub.

[2] Dell’Acqua, F., McFowland III, E., Mollick, E. R., et al. “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality.” Harvard Business School Working Paper 24-013, 2023.

[3] Brynjolfsson, E., Li, D., and Raymond, L. R. “Generative AI at Work.” Quarterly Journal of Economics, vol. 140, no. 2, 2025 (NBER Working Paper 31161, 2023).

[4] Noy, S., and Zhang, W. “Experimental evidence on the productivity effects of generative artificial intelligence.” Science, vol. 381, no. 6654, 2023, pp. 187–192.

[5] Juran, J. M. Quality Control Handbook. McGraw-Hill, 1951. Operational source of the 80/20 rule, after Pareto (1896).

[6] Anthropic. “Updates to our Consumer Terms and Privacy Policy.” Anthropic, 2025–2026. Training opt-out and retention terms reviewed with the Fieldpiece contract team.

[7] lateralworks FTTM methodology. Synthesis principle, core-team scrub process, and tooling (milestone timeline, milestone matrix, decision and cost-of-delay models). Internal practice documentation, 2026.