



Whitepaper

Portfolio prioritization & starts control

How to govern what enters the active pipeline so the work that starts actually finishes

FTTM Best Practices Series. A lateralworks methodology paper on portfolio prioritization and starts control: AHP-based ranking without constraints, then resource reality applied to determine what actually starts. Connects the discipline to four decades of portfolio research from Saaty, Cooper, Reinertsen, Goldratt, Wheelwright and Clark, and Rogers and Blenko.

Prepared by
lateralworks
FTTM methodology

Date
May 2026
Methodology paper

Online
lateralworks.com
FTTM Best Practices Series

Table of contents

Portfolio prioritization & starts control

Abstract	03
01 Why projects start late	04
02 The portfolio decision-making flow	07
03 The AHP-based prioritization model	10
04 Starts control: the pipeline gate	14
05 Decision rights inside the portfolio	18
06 Outside research that anchors the discipline	22
07 Implementation: making it real	25
08 Conclusion	29
References	32
Appendix A — Altera (programmable logic)	34
Appendix B — SunPower (solar cells)	39

Core thesis. Most product organizations start more projects than they can finish. The cure is not better execution. It is a portfolio discipline that ranks every candidate against weighted strategic and economic criteria, then gates entry to the active pipeline against real resource capacity. Rank without constraints first. Apply constraints second. One person owns the call. Starting fewer projects is how more projects finish on time.

Overview

Abstract

The single most damaging failure in product development portfolio management is starting more projects than the organization can execute. The symptom is familiar: every project is a priority, every team is fully committed, and most milestones slip. The root cause is structural. No filter prevents low-value work from entering the pipeline, no mechanism orders what does enter against real resource capacity, and no single person owns the call. Each unit of unfunded work that still gets started consumes capacity from work that is already in flight.

This paper describes the lateralworks FTTM portfolio discipline: a structured prioritization model based on the Analytic Hierarchy Process [1] paired with explicit starts control at the entry to the active pipeline. The flow is sequential. Strategic and economic filters rank every candidate without constraints. A six-step zero-based budgeting (ZBB) loop builds a baseline ranking, captures stakeholder weighting through pairwise comparisons, consolidates consensus, and only then applies budget and headcount limits. The projects above the resource line enter the active pipeline; the rest are deferred or killed.

The discipline draws on independent research streams that converge on the same conclusion. Saaty's AHP [1] formalized the math of multi-criteria decision-making. Cooper, Edgett, and Kleinschmidt documented portfolio management as the highest-leverage practice distinguishing best-in-class new product development [2, 3]. Reinertsen showed why running engineering at high utilization makes queue times explode [4]. Goldratt's Theory of Constraints reached the same conclusion from operations research: limit work-in-progress, finish what you start [5]. Wheelwright and Clark made the case at Harvard Business School for portfolio-level aggregation as the lever on cycle time [6]. Rogers and Blenko's RAPID work [7] supplies the decision-rights model that prevents governance from collapsing into committee.

For semiconductor and systems-product organizations, the prescription is operational. Build the unconstrained ranking first so strategic and economic priorities are not distorted by whoever's loudest about resource pain. Apply the constraint model second, transparently, with the line between funded and deferred visible to the organization. Name a single Decider. Run a quarterly review cadence with explicit emergency-review triggers. The organizations that adopt this discipline ship more products on original calendar; the ones that don't ship many products late.

Key finding. A perfectly prioritized project list is operationally useless if every project on it enters the pipeline at once. The pipeline gate — starts control — is what turns a ranked list into a feasible plan.

01

Section one

Why projects start late

Late delivery is a systems problem, not a person problem. Eight recurring causes show up across organizations regardless of industry, size, or maturity. Every one of them is structural — embedded in how the portfolio is governed, how starts are decided, and how resources are allocated. Section one names the eight causes, then frames the rest of the paper as the structural answer.

Section one

Why projects start late

Ask the leadership team of a product organization why projects finish late and the typical answers are individual: a particular engineer left, a supplier missed a date, a customer changed scope. Aggregate the same answers across many programs over many years and a different picture emerges. The same eight causes appear over and over, and they are structural rather than personal. They describe how work enters the system, how it is prioritized, and how resources are matched to it.

Project Management Institute research is consistent on this point. PMI's Pulse of the Profession surveys repeatedly find that organizations with mature portfolio management complete significantly more projects on time and on budget than those without [8]. The differentiator is not execution skill. It is the discipline that decides which projects enter the pipeline and in what order.

What causes late projects...

- starts late
- resources tied up on last late project
- no priority, all equal
- if prioritized, it changes through development
- no management of FFE
- under resourced (skills & qty)
- over-committed resources
- poor/unmanaged execution
- no single person owns the portfolio

Figure 1. What causes late projects — the recurring structural causes that aggregate across hundreds of programs. Each cause is embedded in how work enters the pipeline rather than in how teams execute it.

The eight causes

Project starts late. The simplest cause is also the most common: the calendar baseline assumed a start date that did not happen. Resources were tied up on the previous late program. The project lost its planned launch window before any technical work began. Late starts compound at every downstream milestone gate.

Resources tied up on the last late one. When the previous program ran past its commit date, the engineers earmarked for the next program are still consumed by the old one. The cascade is mechanical: late finishing N produces late starting N+1, which produces late finishing N+1.

No priority — every project equal. If everything is a priority, nothing is. When the organization cannot articulate which two or three programs matter most, the local optimization that happens at every team and

every gate has no signal to follow. Decisions get made on whichever stakeholder is most insistent.

Priorities shift mid-flight. Even when a priority is declared, it often does not survive the first quarterly review. Reinertsen called this the "fog of priorities" — the cumulative cost of recurring re-prioritization is a hidden tax that does not show up in any program plan but shows up in every program's schedule slip [4].

Fuzzy front end ungoverned. The period between identifying an opportunity and formally starting a project is where most late projects are born. Smith and Reinertsen called this the fuzzy front end and showed that managing it as a focused, time-boxed activity is the single biggest lever on overall cycle time [9].

Under-resourced on skills. Headcount totals can look fine while specific skill gaps are crippling. A program that needs a particular kind of analog designer or RF integration lead can stall for months because the organization counts people but does not count capabilities.

Over-committed on quantity. The mirror of the skill problem. The organization adds projects to the active list faster than it finishes them. Each engineer carries three or four concurrent assignments, switching costs dominate, and individual productivity collapses. The companion lateralworks paper on the optimal project load documents the quantitative penalty in detail [10].

No single owner of the portfolio. The most diagnostic of the eight. When portfolio decisions belong to a committee, the committee optimizes for consensus, not for cycle time. There is no one whose job is to say no, and so no one says no. Rogers and Blenko's research at Bain framed this as the foundational governance failure in slow organizations [7].

Key finding. Late projects are not a sign of weak teams. They are a sign of weak portfolios. The intervention point is upstream of execution: how work is selected, prioritized, and gated into the active pipeline.

02

Section two

The portfolio decision-making flow

The portfolio process operates as a progressive filter. Strategic fit eliminates work that does not belong in the pipeline at all. Economic test ranks the survivors by expected value. Constraint modeling translates the ranked list into a feasible set of starts. Failure analysis closes the loop, feeding what was learned back into the next round. Section two walks the flow from end to end.

Section two

The portfolio decision-making flow

The FTTM portfolio process answers four questions in sequence. Where should we focus? What should we start? In what order, given real constraints? What did we learn? Each question filters the work further. Each stage eliminates options that do not belong, so the organization does not waste planning effort on projects that will never make the cut anyway.

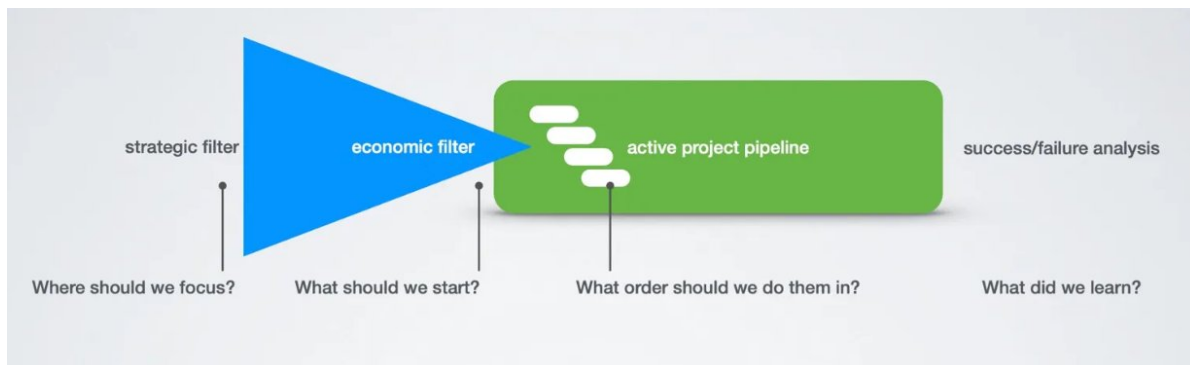


Figure 2. The portfolio decision-making flow. Strategic filter (where should we focus?), economic filter (what should we start?), active pipeline (in what order under which constraints?), and success/failure analysis (what did we learn?) form the four-stage progressive filter that converts candidate ideas into shipped products.

Strategic filter — where should we focus?

The strategic filter establishes the boundaries of the portfolio. Which market segments, technology domains, and customer problems align with the long-term direction? Projects that fail this filter are eliminated before any economic analysis begins. This prevents the common failure mode where the organization funds work that looks economically attractive in the short term but does not belong in the long-term strategy.

Cooper, Edgett, and Kleinschmidt's benchmark research at McMaster University documented this point repeatedly: the best-performing innovators are the ones with explicit strategic buckets and explicit rules about what gets in [2, 3]. The bucket definitions are written down. The rules are followed. Programs that do not fit — however attractive on their own merits — are not funded.

Economic filter — what should we start?

Projects that survive the strategic filter are then ranked against economic criteria. Revenue potential. Strategic fit. Market-maker engagement. Differentiation versus competition. Execution risk. Each criterion gets a weight, and projects are scored against the weighted criteria. The output is a ranked list ordered by composite value. This is the work that the AHP-based model in section three does formally.

Active pipeline — in what order, given constraints?

The transition from ranked list to active pipeline is where starts control operates. A perfectly ranked list is operationally useless if the organization tries to execute every project on it at once. The pipeline stage applies real constraints — engineering headcount, capital equipment, supplier capacity, dependency on predecessor work — to determine which projects actually start, when they start, and which must wait.

Failure analysis — what did we learn?

Completed projects feed back into the strategic and economic models. Lessons on estimation accuracy, market assumptions, and execution challenges improve the fidelity of the next round's decisions. The feedback loop is what prevents the same misjudgments from being repeated. Without it, the model is open-loop and stays whatever it was when it was first calibrated.

03

Section three

The AHP-based prioritization model

The lateralworks portfolio discipline runs on the Analytic Hierarchy Process — a mathematical method developed by Thomas Saaty in the 1970s for converting subjective judgments about relative importance into a rigorous, weighted ranking [1]. The method scales to dozens of criteria and dozens of stakeholders, and it produces a consistency check that flags logically incoherent judgments. Section three walks the six-step process.

Section three

The AHP-based prioritization model

Saaty designed AHP to handle the most common failure of group decision-making: a stakeholder group that knows in principle that criterion A is more important than criterion B, but cannot agree on how much more. AHP turns the abstract question "which criteria matter most?" into a series of concrete pairwise comparisons. For every pair of criteria, a stakeholder makes one judgment: which matters more, and by how much. The math behind AHP guarantees that these binary judgments produce a consistent set of weights, even when the criterion list is large.

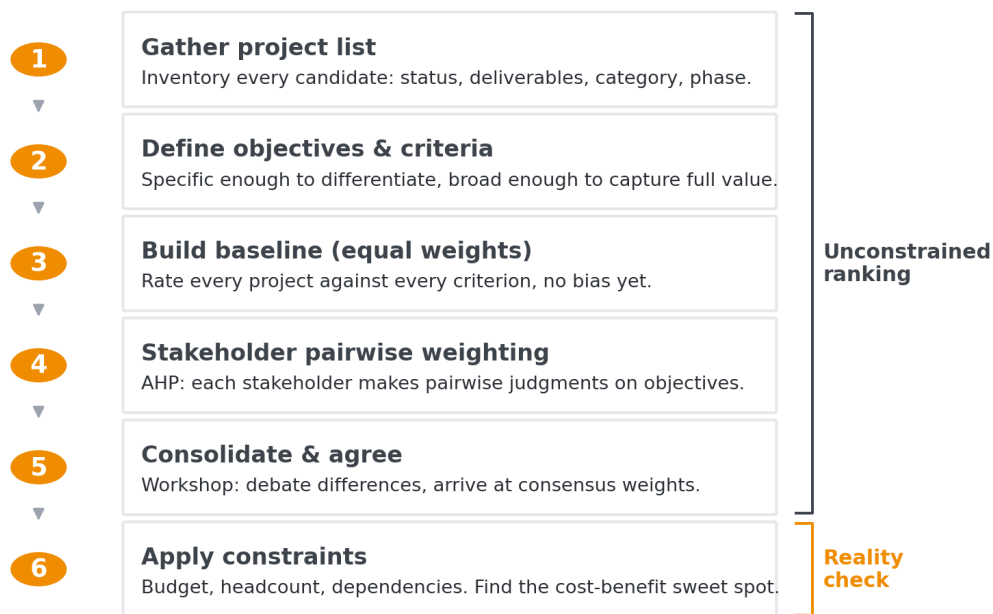


Figure 3. The six-step zero-based budgeting flow. Steps 1–5 produce an unconstrained ranking by strategic and economic value. Step 6 applies the resource reality.

Step 1 — gather the project list

The process begins with a comprehensive inventory. Every candidate project goes on the list: new product initiatives, platform investments, customer-specific developments, technology programs. Each project carries its current status, expected deliverables, category, product series, and phase. The inventory is the input to the prioritization model. Skipping projects that the organization thinks are obvious "must-haves" defeats the purpose — the AHP model is designed to test those assumptions, not accept them.

Step 2 — define objectives and criteria

The criteria must be specific enough to differentiate between projects and broad enough to capture the full value system. A typical lateralworks engagement uses five to seven criteria. Common ones are listed below.

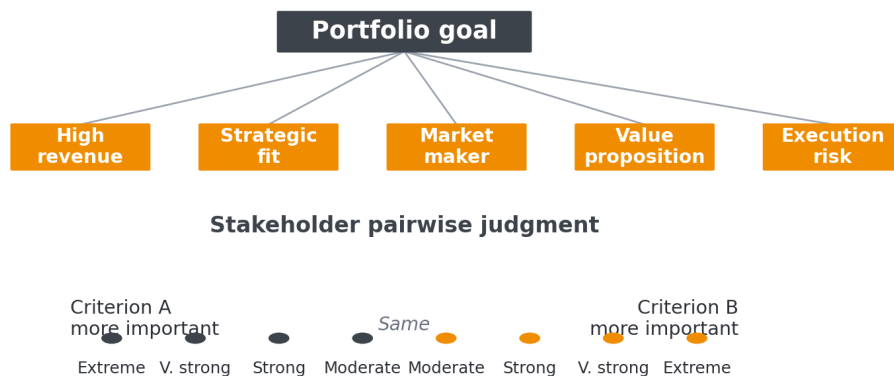
Criterion	Definition	Rating scale
High revenue	Sales-out dollar impact over the planning horizon.	Absolute dollar value estimate.
Strategic fit	Alignment with market segment strategy and technology vision.	Full / Partial / Weak — market and technology sub-criteria.
Market-maker engagement	Opportunity to win designs at target market-maker accounts.	Agreement / High confidence / Uncertain.
Value proposition	Likelihood of winning via differentiation versus competition.	>70% / 50–70% / <50% win probability.
Execution risk	Engineering confidence in schedule and budget.	Low / Medium / High confidence.

Table 1. Example portfolio criteria with definitions and rating scales. The exact set is calibrated per engagement.

Step 3 — build the baseline (equal weights)

Before any stakeholder weighting is applied, every project is rated against every criterion using equal weights — typically one over the criterion count, so 16.7% each for six criteria. The baseline exposes how projects rank on their raw attributes alone, before organizational politics or executive preference influence the output. Surprises at this step are diagnostic. A project everyone assumed was a top priority that ranks in the middle of the baseline is signaling something that the unconstrained model is catching.

Step 4 — stakeholder pairwise weighting



Each pair of criteria gets one judgment. AHP turns the binary calls into a consistent set of weights.

Figure 4. The AHP pairwise comparison. Each pair of criteria gets a single judgment on the 9-point Saaty scale. The math turns the binary calls into consistent weights.

Each stakeholder works through the criterion pairs independently. For each pair, the stakeholder picks which criterion matters more and how much more, on a defined intensity scale: same, moderate, strong, very strong, extreme. Saaty's key insight was that the pairwise judgments are easier to make accurately than direct weight assignments — humans are better at "A versus B" than at "distribute 100 points across these eight items" [1].

AHP also computes a consistency ratio for each stakeholder's judgments. If a stakeholder says A is more important than B, B is more important than C, but C is more important than A, the model flags the inconsistency. The stakeholder is asked to revisit the specific pair that broke transitivity. This is one of the most valuable mechanics of the method: it catches incoherent reasoning before it becomes a portfolio decision.

Step 5 — consolidate and agree on priorities

The consolidation workshop is where the real decision-making happens. Stakeholders come together to review how their individual weightings differ, see what those differences mean for project rankings, and then make new group pairwise judgments. The discussion is structured. Participants argue for their positions. They debate trade-offs. They converge on consensus weights. The final ranking is not imposed; it is negotiated with full transparency into how each criterion's weight affects the outcome.

Common workshop activities include reviewing differences in individual weights and their effect on rankings; making new group pairwise judgments while debating positions; simulating how changing a single criterion weight reshuffles the priority order; filtering rankings by product category or product series to test sub-portfolio behavior; and discussing implications before agreeing on the final ranking.

Step 6 — apply constraints

Steps 1–5 produce the unconstrained ranking. Step 6 applies the resource reality — budget envelope, engineering headcount, shared infrastructure, supplier capacity, dependency on predecessor projects. The model loads projects in priority order until the cumulative resource demand crosses the constraint ceiling. Projects above the line are funded; projects below the line are deferred or killed. This is the work of section four.

04

Section four

Starts control: the pipeline gate

Starts control is the operational mechanism that turns a ranked list into a feasible plan. It is the single most damaging failure point in portfolio management when it is missing, and the single most concrete improvement when it is added. The principle is counterintuitive but well-established: starting fewer projects leads to more projects finishing on time. Section four covers the underlying mechanism.

Section four

Starts control: the pipeline gate

A prioritized project list, no matter how analytically rigorous, is operationally useless if every project on it enters the active pipeline at once. Engineering organizations have finite capacity. When the total resource demand of all active projects exceeds that capacity, every project slows down. Tasks queue for shared resources. Engineers context-switch between programs. Lead times extend. Defect rates rise. By trying to do everything, the organization accomplishes less than if it had focused on fewer initiatives.

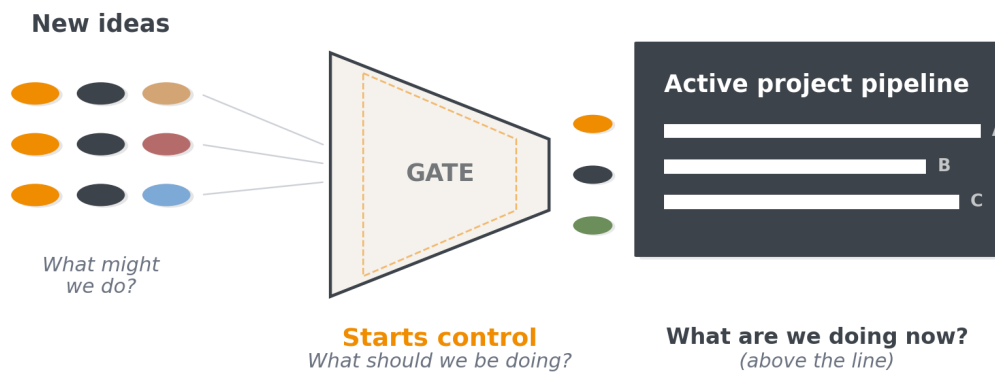


Figure 5. Starts control as the gate between candidate ideas and the active pipeline. Only projects that pass the gate consume engineering capacity.

The resource capacity problem

The mechanism is well-documented in queueing theory. Reinertsen applied the math of M/M/1 queueing systems to product development and showed that wait time grows nonlinearly with utilization [4]. At 60% utilization, queue time is short and predictable. At 90%, it has multiplied roughly nine-fold. At 95%, it spirals. The curve is exponential, not linear — small overcommitment produces large schedule damage.

Goldratt reached the same conclusion from a different starting point [5]. His Theory of Constraints framed multitasking as the silent throughput killer in any system where work units are interdependent. His prescription — limit work-in-progress, focus on the bottleneck, finish what you start — is the operational complement to Reinertsen's quantification of why. The two frameworks reach the same place: accelerate complex work by reducing the number of things in flight, not by increasing them.

Modeling constraints — budget and headcount

Step 6 of the ZBB process translates the unconstrained ranking into a feasible portfolio. Two constraint forms cover most cases. The simplest is total R&D budget per project, expressed in dollars over the planning window. The more granular is engineering headcount per project per quarter, which captures the skill-mix reality the dollar number obscures.

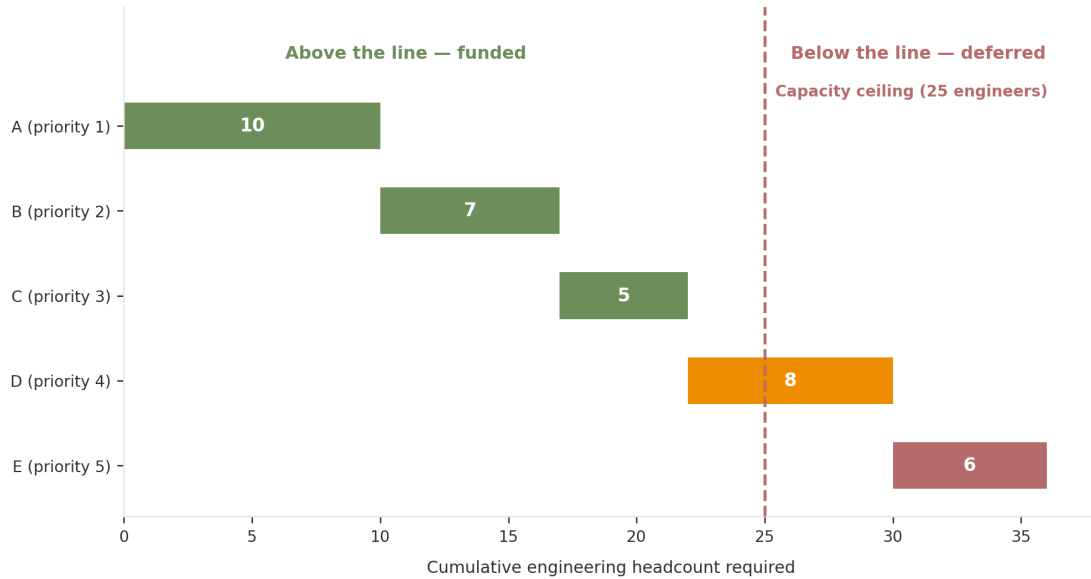


Figure 6. Resource capacity assessment with cumulative headcount plotted against the available capacity ceiling. The line between funded (above) and deferred (below) is explicit.

Projects are loaded in priority order. As each project is added, its resource demand accumulates. The point at which cumulative demand crosses the constraint ceiling defines the line. Projects above the line are funded and staffed. Projects below the line are deferred, descoped, or killed. The exercise is not static. The constraint model is re-evaluated whenever priorities shift, projects complete, or capacity changes.

The cost-benefit sweet spot

Cost-benefit analysis at the portfolio level reveals an effect that is intuitive once seen but rarely acted on: portfolio benefit does not scale linearly with investment. The first dollars deployed fund the highest-priority projects and deliver disproportionate value. As lower-priority projects get added, marginal return diminishes. At some point, the next project delivers negligible incremental benefit while consuming significant resources and adding management overhead.

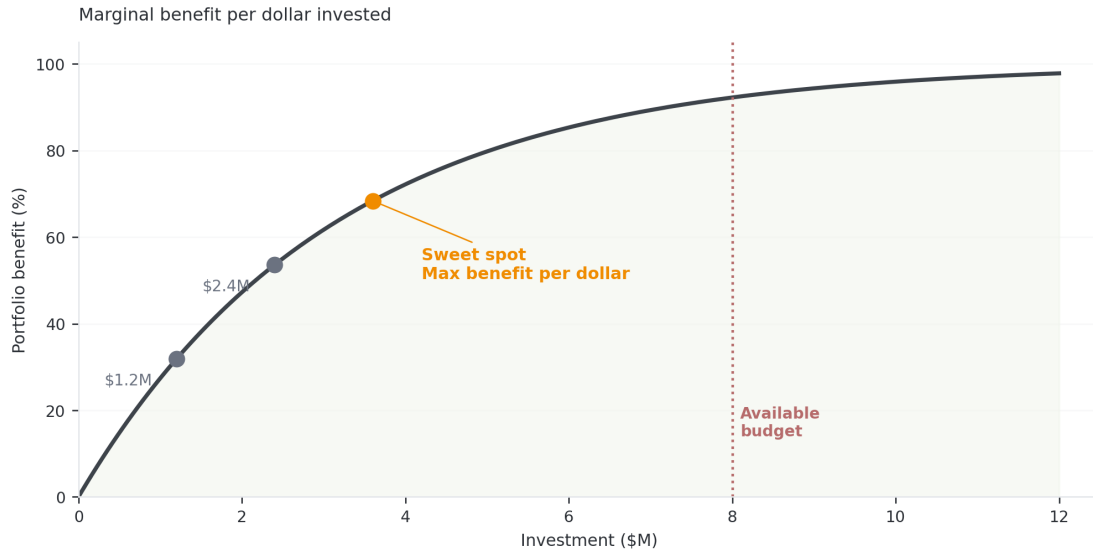


Figure 7. The portfolio sweet spot. Below the spot, the organization is under-investing and leaving value on the table. Above it, additional investment yields diminishing returns and risks overloading execution capacity.

The sweet spot is the investment level where the benefit curve begins to flatten. Cooper, Edgett, and Kleinschmidt's benchmarking data confirms this empirically across hundreds of companies: the best-performing portfolios are not the largest ones, they are the ones tuned to the marginal-benefit point [3]. The optimal portfolio is not the one that funds every good idea. It is the one that maximizes total delivered value within realistic constraints.

Determining when projects start

With the constrained portfolio defined, starts control governs the timing of project initiation. The decision inputs are explicit: the project's position in the priority ranking; availability of the specific engineering skills and infrastructure required; current utilization of shared resources; status of any predecessor projects or technology dependencies; and the business case including the cost of delay if the start is pushed.

When a project completes or is killed, the freed capacity flows to the next project in the priority queue. This pull-based system keeps the pipeline loaded to capacity but never overloaded. It also creates natural urgency: teams know that finishing their current program faster directly enables the next priority to start, which most teams find motivating once the system is visible to them.

Key finding. Starting fewer projects is the highest-leverage lever on cycle time available to the organization. The mechanism is mechanical — queue length governs flow — and the data has been documented across software, hardware, manufacturing, and service organizations for four decades.

05

Section five

Decision rights inside the portfolio

A portfolio process is only as effective as the governance that sits behind it. The most analytically rigorous AHP model produces no decisions if no one is accountable for making the call. Section five covers the decision-rights model that lateralworks uses to prevent governance collapse: the IRADP framework, adapted from Rogers and Blenko's RAPID work at Bain.

Section five

Decision rights inside the portfolio

The fundamental governance question for any portfolio is who decides what enters the active pipeline. The default answer in most organizations is "the leadership team" or "the executive staff." Both formulations describe a committee, and committees optimize for consensus rather than for cycle time. When a committee owns the portfolio, the path of least resistance is to add another project to the active list rather than to say no to the executive whose project would be deferred.

Rogers and Blenko's research at Bain showed that organizational speed correlates strongly with decision-rights clarity, and that unclear rights account for the majority of slow strategic decisions in their dataset [7]. Their RAPID framework names five roles for any decision: Recommend, Agree, Perform, Input, and Decide. lateralworks uses an adapted variant called IRADP, which reorders the roles into the explanatory hierarchy below.

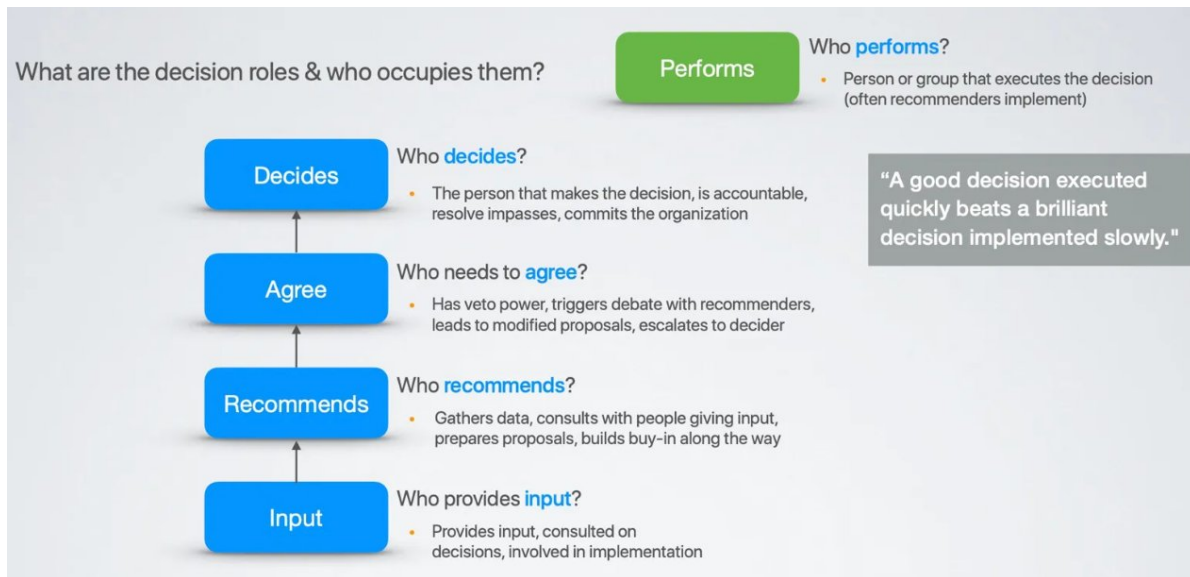


Figure 8. What are the decision roles and who occupies them. Each role is named, each role is assigned to a specific person, and the Decider is singular. "A good decision executed quickly beats a brilliant decision implemented slowly."

Who occupies each role

Input is provided by anyone with relevant data or perspective. Functional managers contribute information on resource availability and skill mix. Customer-facing teams contribute voice-of-customer signal. Engineering leads contribute feasibility judgments. Input is consultation, not authority — the input role does not trigger debate or escalate decisions.

Recommends is the role that does the actual analytical work. The recommender gathers data from the input network, builds proposals, runs the AHP model, and prepares the constraint analysis. The recommender also builds buy-in along the way — walking through the model with the people whose work it affects so the consolidation workshop is a refinement rather than a first exposure.

Agrees has veto power. The agree role triggers debate when a recommendation conflicts with a constraint that the recommender did not see. In a portfolio process, this is typically a functional VP whose organization would be asked to deliver beyond its capacity, or a finance leader whose budget envelope would break. Veto leads to a modified proposal or escalation to the Decider — not to indefinite stalemate.

Decides is the role that makes the call. The Decider is singular. They resolve impasses, commit the organization, and are accountable for the outcome. In a portfolio context, the Decider is most often the general manager, division head, or product line leader who owns the P&L. In a startup, it is the CEO. The critical design choice is that the role is owned by one person. Without a single Decider, decisions devolve into consensus-seeking that avoids the hard trade-offs.

Performs is the role that executes the decision. In portfolio governance, the recommenders are typically also the performers — they run the workshop, update the model, and communicate the outcome. The roles are different even when the people are the same, and naming them separately preserves the distinction.

Why a single Decider matters

The most diagnostic governance pathology in slow organizations is a portfolio that no one owns. When ownership is plural, every individual stakeholder is incentivized to protect their own projects, and the equilibrium is a bloated active list with no one accountable for the cycle-time outcome. A single Decider — with the authority to say no, the visibility into the trade-offs, and the accountability for the calendar — is the structural mechanism that breaks the equilibrium.

In the IRADP context, the portfolio owner is the Decider for what enters the pipeline. Functional managers provide input on resource availability. Product managers and recommenders run the AHP model and propose start orderings. Engineering and finance leadership must agree the proposal is feasible. But one person decides, and the decision is binding for the planning horizon.

Key finding. Governance without accountability produces no decisions. The organizational performance lever is naming a single Decider for the portfolio, with explicit authority and explicit cycle-time accountability.

On focus

The constraint that matters

The most important variable in product development is queue length, not utilization or efficiency.

— Donald G. Reinertsen

The Principles of Product Development Flow (2009)

06

Section six

Outside research that anchors the discipline

The lateralworks portfolio discipline is not an opinion. It is the operational synthesis of four decades of independent research streams that converge on the same conclusions: rank without constraints first, gate against capacity second, name a single Decider, do fewer things at once. Section six walks the streams and shows how each one supports a specific element of the FTTM practice.

Section six

Outside research that anchors the discipline

Six research streams supply the empirical foundation for the portfolio discipline described in sections two through five. Each stream developed independently, in a different field, and each reaches the same conclusion from a different starting point. The convergence across streams is the strongest argument for the discipline: it is not a theory from one school but a pattern that shows up wherever complex work is governed against finite capacity.

Saaty and the Analytic Hierarchy Process

Thomas Saaty developed AHP at the Wharton School in the 1970s to solve a class of decision problems that classical optimization could not handle: multi-criteria decisions where the criteria themselves had different weights, and where the weights were subjective rather than objective [1]. His insight was that humans make pairwise comparisons more reliably than they make absolute rankings, and that pairwise comparisons can be combined into a consistent eigenvector through a well-defined matrix operation. AHP has since been applied across thousands of domains — capital allocation, supplier selection, public policy, medical diagnosis. The portfolio prioritization use case is a direct application of Saaty's original framework.

Cooper, Edgett, and Kleinschmidt on portfolio management

Robert Cooper's Stage-Gate process [11] is the most widely adopted innovation framework in industry, and Cooper's subsequent benchmarking work with Edgett and Kleinschmidt at McMaster University documented portfolio management as the single highest-leverage practice that distinguishes best-in-class innovators from the rest [2, 3]. Their data, drawn from hundreds of firms across two decades, shows that the firms with explicit portfolio processes — strategic buckets, scoring models, go-kill discipline — ship significantly more new products on time and at higher hit rates than firms without. Their research is the empirical case for treating portfolio governance as a first-order strategic capability rather than as an administrative function.

Reinertsen on queue length and cost of delay

Donald Reinertsen's Principles of Product Development Flow [4] is the source for the queueing-theory analysis of starts control. Reinertsen showed that product development behaves mathematically like a telecommunications network: introduce variability and increase utilization, and queue times grow exponentially. His second contribution was the cost-of-delay framing — the right performance metric is not utilization but the revenue impact of each week of schedule slip. Both insights are foundational for starts control: the queue math explains why overloading creates so much schedule damage, and the cost-of-delay frame explains why leaving capacity unused is the rational choice.

Goldratt and the Theory of Constraints

Eliyahu Goldratt's Theory of Constraints, articulated in The Goal and extended in Critical Chain [5], reaches the same conclusions as Reinertsen from a different methodological starting point. Goldratt's focus was the manufacturing shop floor, where multitasking on machine tools and shared resources destroyed throughput. His prescription — limit work-in-progress, identify the system's constraint, subordinate everything else to the constraint — maps directly onto product development portfolio management. The starts control mechanism is the WIP limit. The shared engineering capacity is the constraint. The portfolio discipline subordinates everything else to keeping that constraint productive.

Wheelwright and Clark on the aggregate project plan

Steven Wheelwright and Kim Clark's Revolutionizing Product Development [6] introduced the aggregate project plan: an organization-wide view of all active development projects, classified by type, sized by resource demand, and sequenced deliberately rather than accreted by accident. Their case studies across automotive, semiconductor, and consumer electronics firms documented 30 to 50 percent cycle-time reductions in organizations that adopted the aggregate-plan discipline. The lateralworks portfolio process is an operational implementation of the aggregate project plan with the AHP scoring layer added.

Rogers and Blenko on decision rights

Paul Rogers and Marcia Blenko's 2006 Harvard Business Review article "Who Has the D?" [7] is the source of the RAPID framework that lateralworks adapts as IRADP. Their research at Bain showed that decision-rights clarity correlates strongly with both organizational speed and execution quality, and that ambiguous rights are a leading explanation for slow strategic decisions. The single-Decider principle in section five is a direct application of their finding that without one accountable person, decisions devolve into committee dynamics.

Convergence across streams

The streams are independent. Saaty came from operations research at the University of Pennsylvania. Cooper's work was done at McMaster University in Ontario. Reinertsen built his framework as an industry consultant. Goldratt was a physicist who turned to manufacturing. Wheelwright and Clark were Harvard Business School researchers. Rogers and Blenko worked at Bain. They did not coordinate, they did not cite each other (in most cases), and they reached the same conclusions about how complex work is governed: rank candidates rigorously, limit work in progress, name single accountability, and subordinate everything else to finishing what you started.

Key finding. The portfolio discipline is not derived from a single school. It is the operational pattern that emerges when multiple independent research streams are read together, and it has held up across forty years of empirical validation in software, hardware, manufacturing, and services.

07

Section seven Implementation: making it real

Implementing the portfolio discipline is not a one-time event. It is a sustained organizational change that runs on a regular cadence, with explicit triggers for unscheduled review and explicit feedback loops to keep the model honest. Section seven walks the implementation sequence and the operating cadence that has worked across lateralworks engagements.

Section seven

Implementation: making it real



Figure 9. The portfolio modeling flow. The model without constraints ranks projects by strategic fit and economic test (projects below the line are eliminated). The model with constraints prioritizes starts against the reality of resources, budgets, risks, and dependencies. Execution refines the model with time-phased information; the loop iterates.

Step one: stand up the model

The first step is to populate the AHP model for a single product line or division. A lateralworks engagement typically runs this as a four-week sprint. Week one: gather the project list. Week two: define criteria and rating scales. Week three: stakeholders make individual pairwise judgments. Week four: consolidation workshop. The output is a constrained portfolio with named owners and a visible line between funded and deferred projects.

The decisive choice in the first sprint is who facilitates the workshop. The right facilitator is neutral on the outcome — not invested in any specific project surviving the constraint cut. In most engagements this is an outside consultant or a corporate strategy lead with no operational stake. The wrong facilitator is a functional VP whose own programs are on the candidate list. They cannot help advocating, and the workshop loses its analytical character.

Step two: implement starts control governance

With the model standing, the governance structure is the next piece. Assign the portfolio Decider — one person, named in writing, with explicit authority over what enters the pipeline. Establish the review cadence: quarterly for strategic and economic reassessment, monthly for portfolio health checks, weekly for execution-level refresh planning. Define the emergency review triggers in advance.

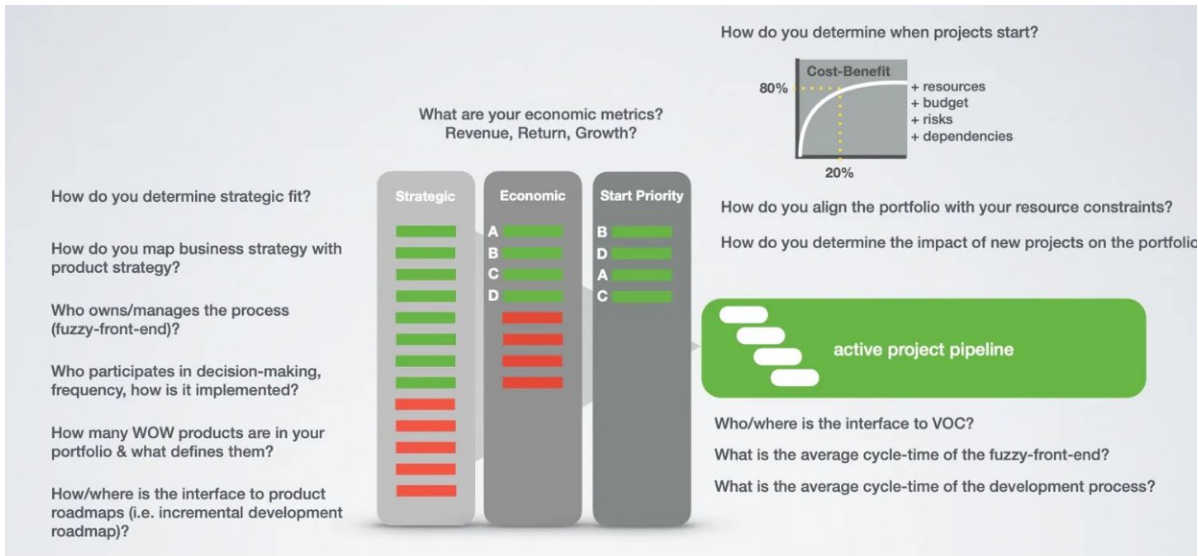


Figure 10. Who owns the portfolio. The Decider runs the diagnostic on the left (strategic fit, business-product mapping, fuzzy front end governance, decision-making process, WOW product count, roadmap interface), the constraint diagnostic on the right (resource alignment, impact of new starts, cycle-time of front-end and development), and the cost-benefit sweet spot above. The strategic, economic, and start-priority columns make the constraint cut visible to everyone.

Common emergency triggers include loss of a key customer, significant change in available engineering capacity, a major project delay that threatens to consume capacity earmarked for the next priority, and a competitive event that changes the cost-of-delay calculus. The triggers should be defined during the requirements and feasibility phases of each program so the threshold for escalation is objective rather than emotional.

Step three: integrate with the broader system

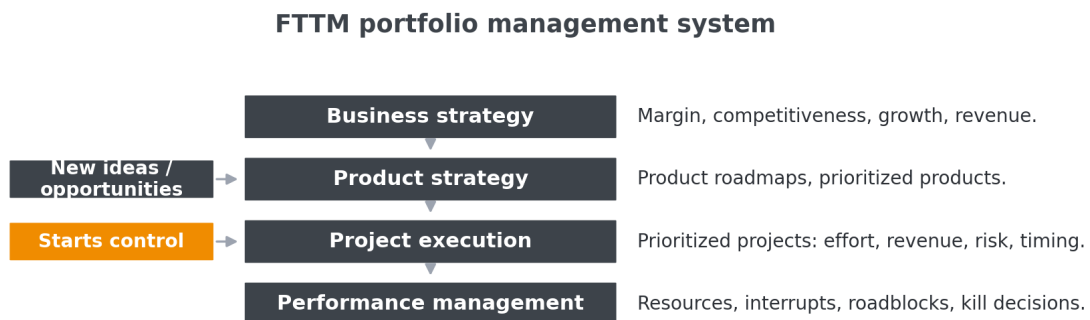


Figure 11. The complete FTTM portfolio management system. Business strategy and product strategy feed the model; new ideas and starts control gate the pipeline; execution and performance management close the loop.

The portfolio process does not stand alone. It connects upstream to business strategy (which defines the strategic buckets), parallel to product strategy (which translates buckets into roadmaps), and downstream to

project execution and performance management. The connections matter as much as the model itself. A portfolio that produces good rankings but is disconnected from how the organization actually executes — because performance management still tracks utilization rather than cost of delay, or because product strategy is generating new opportunities faster than the gate can absorb them — will not deliver the cycle-time benefit that the analytics suggest.

Step four: iterate and improve

Hold quarterly reviews of projects, phases, and priorities for each segment. Include a feedback session on the process and tooling. Refine criteria weights as the organization learns which criteria actually predict project success and which were directional but not decisive. The model is not static. It is a living representation of how the organization makes its largest capital-allocation decisions, and it should improve every cycle as the data quality improves.

Common implementation pitfalls

Three failure modes recur often enough to be worth naming. **The first** is using AHP to ratify a predetermined answer. A leadership team that has already decided which two projects they want funded, and is running the model to validate the decision, will arrange the criteria weights to produce that output. The model is a transparency tool, not an arbitration tool, and it loses its value when used for the latter.

The second is letting the constraint analysis become private. When the line between funded and deferred is hidden, the projects below the line continue to consume capacity informally — a half-headcount here, a side analysis there. The constraint discipline only works when the line is visible to everyone, including the engineers whose programs sit just below it.

The third is letting the Decider role drift back to a committee. Within six months of a successful first cycle, the organization's default committee dynamics tend to reassert themselves. The portfolio Decider needs explicit reinforcement from the executive layer above — visible meetings, written authority, and consistency in escalation paths — or the role erodes back into consensus.

08

Section eight

Conclusion

The portfolio discipline solves one problem: ensuring that the right projects start at the right time, given the reality of finite resources. The mechanics are simple. The discipline to sustain them is what separates organizations that ship on calendar from organizations that do not.

Section eight

Conclusion

Six principles run through this paper. Each is structural, codifiable as policy, and observable in operating practice.

Principle	Application
Rank without constraints first	Use AHP pairwise comparisons to rank projects by strategic and economic value before applying resource limits. This prevents constraint-driven thinking from distorting the priority signal.
Apply constraints second	Model budget and headcount limits to determine what can actually be funded. The line between funded and deferred is explicit, transparent, and driven by data.
Find the sweet spot	Use cost-benefit analysis to identify where marginal returns diminish. Do not fund every good idea. Fund the portfolio that maximizes total delivered value.
Gate pipeline entry	Starts control prevents overloading. Projects enter the active pipeline only when capacity exists and priority warrants it. Starting fewer projects means more finish on time.
Assign a single Decider	One person owns the portfolio. They resolve impasses, make trade-offs, and commit the organization. Governance without accountability produces no decisions.
Iterate continuously	The portfolio model is not a static plan. Regular review cadences, emergency triggers, and feedback loops keep the portfolio aligned with reality.

Table 2. Six core principles of FTTM portfolio prioritization and starts control.

The methodology is tool-supported but not tool-dependent. lateralworks uses an Excel-based AHP implementation called fastDecisionAI that any product organization can adopt. The principles above — transparent prioritization, explicit constraint modeling, disciplined starts control, clear decision governance — apply to any product development organization that needs to do more with less by doing less at once.

For semiconductor and systems-product organizations operating against external customer dates and competitive technology windows, the prescription is operational. Stand up the AHP model in a focused four-week sprint. Name the Decider. Make the constraint line visible. Run the cadence. The cycle-time benefit shows up within two quarters — not as a theoretical projection, but as an observable shift in the proportion of programs that hit their original dates.

The discipline is the deliverable. AHP is a method. Starts control is a mechanism. The IRADP framework is a structure. The organizational benefit comes from running all three together, on a cadence, with a single Decider, until the practice becomes the default mode of how the portfolio operates.

R

Sources

References

The references below are the primary sources for the empirical and methodological claims in this paper. Citations are numbered in order of first appearance in the body. Reference 10 cross-links to the companion lateralworks paper on the optimal project load; reference 13 represents lateralworks engagement data across two decades of programs.

Sources

References

- [1] Saaty, T. L. *The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation*. McGraw-Hill, 1980. Republished by RWS Publications, 1990.
- [2] Cooper, R. G., Edgett, S. J., and Kleinschmidt, E. J. "Best Practices for Managing R&D Portfolios." *Research-Technology Management*, 41(4), 1998, pp. 20–33.
- [3] Cooper, R. G., Edgett, S. J., and Kleinschmidt, E. J. *Portfolio Management for New Products*. 2nd edition, Basic Books / Perseus, 2001.
- [4] Reinertsen, D. G. *The Principles of Product Development Flow: Second Generation Lean Product Development*. Celeritas Publishing, 2009.
- [5] Goldratt, E. M. *Critical Chain*. North River Press, 1997.
- [6] Wheelwright, S. C., and Clark, K. B. *Revolutionizing Product Development: Quantum Leaps in Speed, Efficiency, and Quality*. The Free Press, 1992.
- [7] Rogers, P., and Blenko, M. "Who Has the D? How Clear Decision Roles Enhance Organizational Performance." *Harvard Business Review*, January 2006, pp. 53–61.
- [8] Project Management Institute. *Pulse of the Profession: The High Cost of Low Performance*. PMI, 2014, and annual updates 2015–2024.
- [9] Smith, P. G., and Reinertsen, D. G. *Developing Products in Half the Time: New Rules, New Tools*. 2nd edition, John Wiley & Sons, 1998.
- [10] lateralworks. "The optimal project load: how dedicated resources accelerate new product development." Methodology paper, FTTM Best Practices Series, May 2026.
- [11] Cooper, R. G. "Stage-Gate Systems: A New Tool for Managing New Products." *Business Horizons*, 33(3), 1990, pp. 44–54.
- [12] Brooks, F. P. *The Mythical Man-Month: Essays on Software Engineering*. Addison-Wesley, 1975 (anniversary edition 1995).
- [13] lateralworks. "Internal engagement database." Field observations across client programs in semiconductors, systems, medical devices, and industrial products, 2010–2026.
- [14] McKinsey & Company. "What Makes Product Teams Effective?" *McKinsey Digital*, 2024.
- [15] Isaacson, W. *Steve Jobs*. Simon & Schuster, 2011.

A

Appendix A

Altera — programmable logic

Altera Corporation, the programmable logic semiconductor company acquired by Intel in 2015, ran a multi-segment portfolio across industrial, automotive, military, communications, and computer/storage markets. Altera adopted the FTTM portfolio discipline at division level under the name "ABC Division Planning Process" and used it to govern the flow of FPGA, IP, and platform development programs through a sequence of phases from strategy to launch. The discipline survived the Intel acquisition and the patterns documented here are representative of how the methodology lands inside a multi-billion-dollar semiconductor business unit.

This appendix walks the implementation. The strategic and economic filters from section two of this paper sit at the top of the flow. The AHP-based prioritization model from section three lands as the "fastDecisionAI" tool used by every product line manager. The starts control mechanism from section four operates as a sequence of phase gates — strategic, requirements, feasibility, execution, go-to-market — with explicit emergency-review triggers defined at each gate. The decision rights from section five map onto the named roles in the flow: PLM as the recommender and downstream phase owner, engineering manager as agreeer at the requirements-to-feasibility gate, and the executive staff as Decider on segment-level budget.

ABC Division Planning Process - Phases

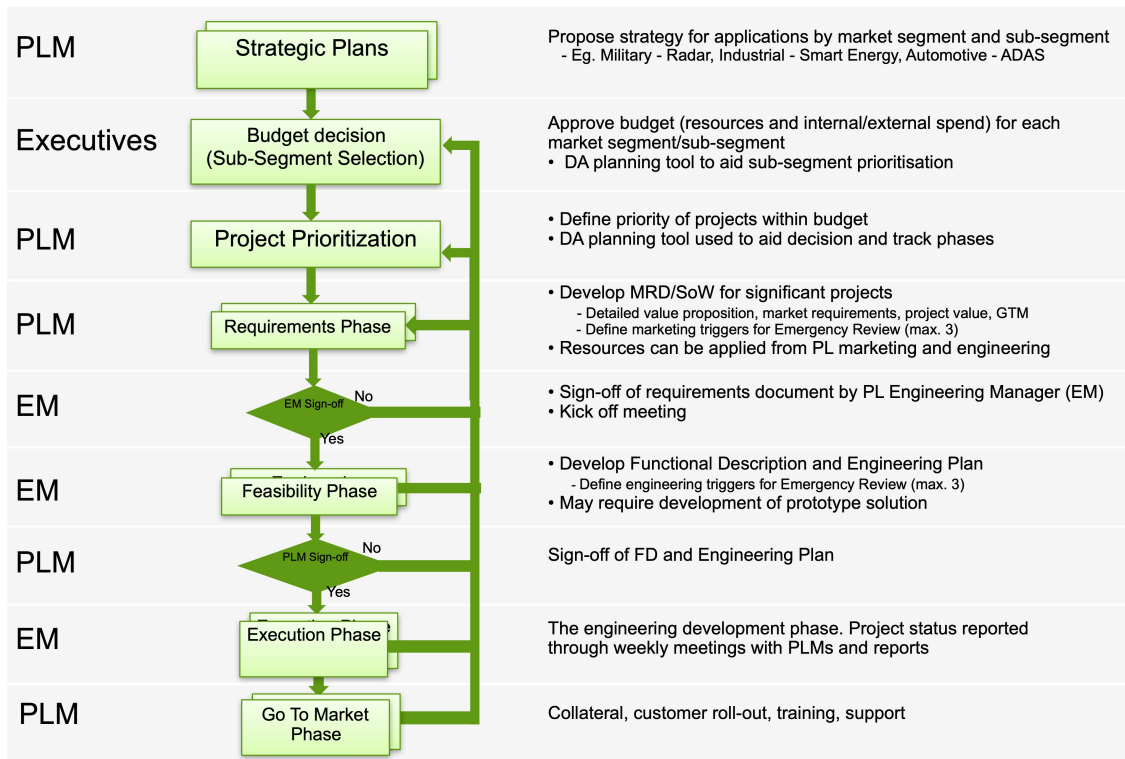


Figure A1. The Altera ABC Division Planning Process: a nine-stage sequence from strategic plans to go-to-market, with named role ownership (PLM, Executives, Engineering Manager) at every stage.

Section A.1

How the discipline maps to the FTTM model

The Altera flow lines up with the FTTM portfolio discipline at every level. The first three stages — strategic plans, budget decision, and project prioritization — are exactly the strategic and economic filters described in section two. The strategic plan stage establishes which market segments and sub-segments belong in the portfolio at all (smart energy, ADAS, industrial communications, radar). The budget decision stage allocates resource and external spend envelopes across the segments. Project prioritization within each segment is then the unconstrained AHP ranking that section three describes — and Altera used the fastDecisionAI tool to do it.

The downstream stages — requirements, feasibility, execution, and go-to-market — are the active pipeline. Once a project clears the project-prioritization gate, it consumes engineering and marketing capacity and proceeds through milestone reviews until release. The sign-off gates between phases function as starts-control checkpoints within the active pipeline: requirements signed off by the engineering manager before kickoff, functional description and engineering plan signed off by the PLM before development begins.

Section A.2

Five weighted criteria with sub-criteria

Criteria Definitions and Ratings		
Criteria	Definition	Rating
High Revenue	The Sales Out \$ impact estimate from 2014-2019	\$ value estimate
Strong Strategic Fit	Assesses how well aligned the project is with both the market segment strategy, and XXX's device and technology vision/strategies.	Full alignment = All sub-criteria below Partial alignment = Meets 3 of sub-criteria Weak alignment = <3 of sub-criteria Sub-criteria 1. Market Alignment 1.1 Well aligned to the Strategic Plan 1.2 Future market segment alignment 2. Technology Alignment 2.1 Aligned with XXX's device & tech vision 2.2 Required by other important opportunities
Engage with market maker	Assesses the opportunity to win one or more designs from the target market maker customer list.	Agreement 1. Requirements understood by XXX 2. Market maker agreed to use XXX 3. Commitment proven through docs (MoU, MRD) High Confidence 1. Requirements understood by XXX 2. Customer has indicated high likelihood 3. Commitment NOT proven through docs Uncertain Current status is NOT Agreement and NOT High Confidence
Strong value proposition	An assessment of likelihood of winning target designs via differentiated solutions Vs competition.	>70% >70% likelihood of winning via differentiated solutions Vs competition 50%-70% 50-70% likelihood of winning via differentiated solutions Vs competition <50% <50% likelihood of winning via differentiated solutions Vs competition
Low execution risk	Assesses the engineering confidence level that the specified feature can be developed within performance, power and system cost budget, and within required delivery schedule.	High Risk Engineering confidence < 40%, when considering time to market requirements, effort estimate, available expertise, complexity of dependencies Medium Risk Engineering confidence 40%->80%, when considering time to market requirements Low Risk Engineering confidence >80%, when considering time to market requirements

Figure A2. Altera's five portfolio criteria with definitions and rating scales: high revenue, strong strategic fit, engage with market maker, strong value proposition, low execution risk. The strategic fit criterion is decomposed into market-alignment and technology-alignment sub-criteria to prevent subjective application.

Two design choices in the Altera criterion set are worth noting. First, the strategic fit criterion is decomposed into market alignment and technology alignment sub-criteria — and the rating scale forces a count of which sub-criteria a project actually meets ("full" requires all of them, "partial" requires three of them, "weak" is below three). This prevents the ambiguity that the abstract phrase "strategic fit" invites. Second, the market-maker criterion has three discrete states — agreement, high confidence, uncertain — each with documented preconditions ("market maker has agreed to use Altera if their requirements are met", "customer commitment is proven through documentation"). This converts what would otherwise be a sales-team opinion into an evidence-based rating.

Key finding. Definitions that include preconditions, not adjectives, are what make AHP criteria operate consistently across stakeholders. "Agreement" with three documentary preconditions is a different kind of category than "strong customer interest" — the first survives a personnel change, the second does not.

Section A.3

AHP weighting with the fastDecisionAI

Altera implemented the AHP model in a Microsoft Excel-based tool called fastDecisionAI. Each PLM owned a workbook tab for their segment (industrial, automotive, military, computer/storage, communications), and the tool walked stakeholders through the pairwise comparison process described in section three. The output was a weighted criteria profile and a ranked list of projects, with a budget envelope and cost-benefit analysis built into the same workbook.

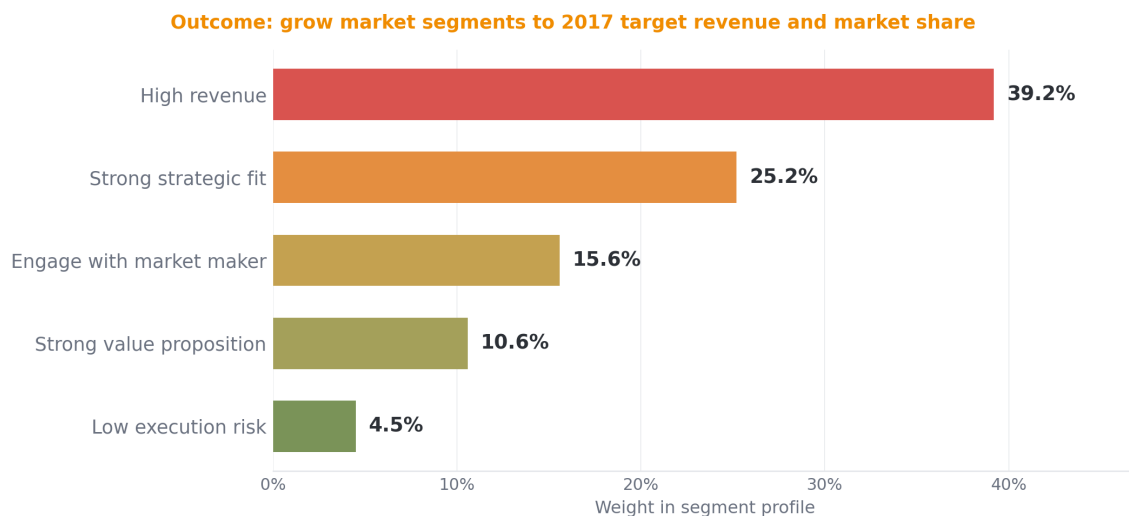


Figure A3. The fastDecisionAI weight profile output for one Altera segment. High revenue carries the largest weight at 39.2%, strategic fit second at 25.2%, market-maker engagement third at 15.6%. Low execution risk receives only 4.5% weight in this segment — execution risk is real but is not a primary lever for priority order.

The weighting profile in Figure A3 is one segment's consensus output. A different segment with different market dynamics produced a different profile — the market-maker criterion typically carried more weight in segments with one or two dominant key-account customers, and less weight in segments with many smaller buyers. The model supports this segment-specific calibration directly: each PLM held their own pairwise judgments in their own workbook tab, and the executive staff compared segments at the budget-decision gate using the model's cross-segment view.

Per-PLM workbook ownership is the operational expression of the IRADP decision-rights model from section five. The PLM is the recommender inside their segment — they gather input from engineering, marketing, and key-account managers, run the AHP weighting workshop, and produce the ranked candidate list. The general manager and executive staff are the agreeers and the Decider at the cross-segment level: they

accept or reject the budget allocations across segments, then trust each PLM's within-segment ranking. This split — within-segment rankings owned by the PLM, cross-segment allocation owned by the executive — keeps the governance load proportional to the decision being made.

Section A.4

Phase tracking and emergency review

The Altera flow runs each project through a sequence of phases — requirements, feasibility, execution, go-to-market — with formal sign-off gates between them. The fastDecisionAI tool tracked phase state alongside the priority ranking, so the executive view of the portfolio showed both ordering (what should we be doing) and state (what stage is each one in). This single-tool combination of strategic and operational view is the reason the discipline survived a decade of Altera leadership changes — the ranking and the pipeline state lived together rather than in separate spreadsheets owned by separate people.

The emergency-review mechanism deserves particular attention. The Altera process required each project to declare up to three marketing triggers and up to three engineering triggers at the requirements and feasibility phases — concrete, objective conditions that, if breached, would force an unscheduled portfolio review. Common triggers were loss of a market-maker customer, change in market conditions that invalidated the original total-addressable-market estimate, significant engineering schedule slip, or loss of key engineering capacity. The triggers were not abstractions — they were named thresholds with defined values, set at the start of the project when judgments were sober rather than under the pressure of a slip already in progress.

Key finding. The right time to define the threshold for an emergency portfolio review is before the project starts. Altera codified up to three marketing and three engineering triggers per program at the requirements and feasibility gates — concrete conditions whose breach would force unscheduled re-prioritization. Defining them early made the escalation objective. Trying to define them in the middle of a slip would have made the escalation political.

The emergency review panel at Altera was not the full executive staff. It was a small group — division senior VP, general manager, and directors most affected by the project — convened on demand. This is the mechanism that prevents the trigger system from collapsing into "another standing meeting." Triggers fire rarely, the panel forms rarely, and when it does form it acts decisively because everyone present has clear authority over the decision being made.

B

Appendix B

SunPower — solar cells

SunPower Corporation, the solar cell and module manufacturer, ran its Cell Technology Development (CTD) organization on a paired discipline: a stage-gate process for product development governance and a zero-based budgeting (ZBB) process for engineering capacity allocation. Together they form the operational implementation of starts control described in section four of this paper. The stage-gate process names the checkpoints; ZBB is what determines which of those checkpoints actually get staffed in any given month.

The case below documents two artifacts. The first is the SunPower stage-gate framework itself, which spans SG0 (ideation) through SG5 (general release) and SG6 (end-of-life). The second is a published CTD ZBB worksheet showing how 23 engineers were allocated across six priority programs over a twelve-month planning horizon — and where, in that horizon, the cumulative resource demand crossed the available capacity ceiling. SunPower called the boundary the "ZBB Line," and visualized it in red across the worksheet. It is the most concrete published example of starts control at work.

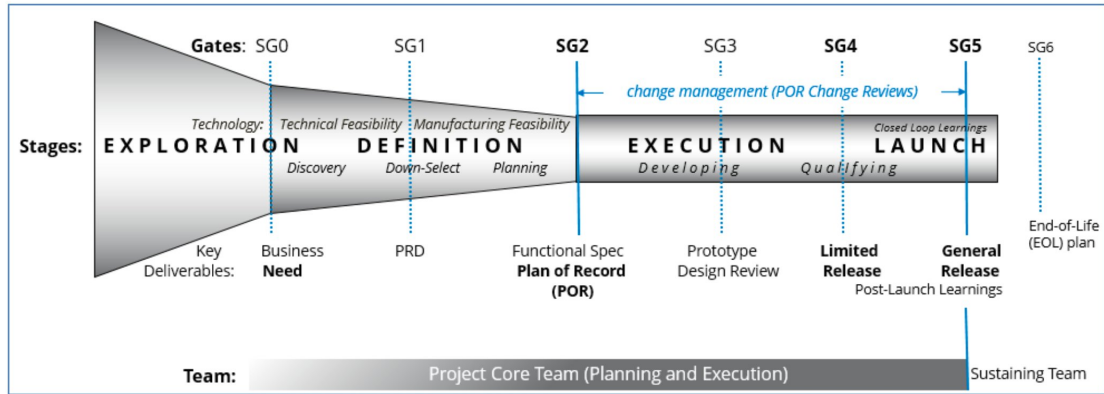


Figure 1. Product Life Cycle Overview

Figure B1. The SunPower stage-gate product life cycle: SG0 exploration through SG5 launch, with key deliverables named at each gate (business need at SG0, plan-of-record at SG2, limited release at SG4, general release at SG5).

Section B.1

A six-stage gate framework

SunPower's stage-gate framework follows the Cooper Stage-Gate model [11] adapted to solar-cell manufacturing. Each gate carries a defined review board, a checklist of expected deliverables, and a decision about whether the program continues, holds, or stops. The framework is flexible by design — small programs run a "SG-lite" subset and large programs run "SG-full" — but the named gates are constant and the governance structure around each gate is the same.



Figure B2. SunPower SG intent: what each gate is meant to achieve, from converging on customer concepts at SG0 to ramping production after SG5. The plan-of-record is set at SG2 — once committed, it is under change control until launch.

Two design choices in the SunPower framework deserve emphasis. The first is the cross-organizational SG approval board: every gate review requires sign-off from nine functional areas (sponsor, customers and stakeholders, development organization, supply chain, upstream manufacturing, downstream operations, finance, quality, reliability) with explicit attendance requirements at each gate. This prevents the common failure mode where engineering signs off on its own readiness and operations later discovers that the manufacturing process is unready.

The second is the change-control discipline beginning at SG2. Once the plan-of-record is set — scope, schedule, and budget all committed in writing — material changes to any of the three trigger a formal POR change review. The change-review process itself defines the threshold between routine adjustment (handled by the project core team) and cross-segment reprioritization (escalated to the SG approval board). This is the same pattern as Altera's emergency-review triggers from Appendix A — but inserted into the active pipeline rather than at the pipeline gate.

Key finding. A stage-gate framework without a defined change-control discipline at plan-of-record degrades into theater within two product cycles. SunPower put the discipline at SG2 — the first gate where scope, schedule, and budget are committed — and required formal POR change reviews for every material adjustment. The result is that the gate decisions made at SG2 mean something six months later, even after the inevitable changes.

Section B.2

The ZBB Line: starts control made visible

SunPower's Cell Technology Development organization had 23 engineers available each month. The CTD planning team translated the stage-gate pipeline into a monthly resource-loading worksheet, with every active and proposed program decomposed into named engineer-month allocations. Programs were loaded in priority order — the unconstrained ranking from the AHP-style filter — and cumulative engineer-months were subtracted from the available pool until the pool went negative. The point at which the pool ran out was called the "ZBB Line." Programs above the line were funded; programs below the line had to be deferred, descope, or stopped.

CTD ZBB worksheet — remaining engineering capacity (23 available)

	Jan'19	Feb'19	Mar'19	Apr'19	May'19	Jun'19	Jul'19	Aug'19	Sep'19	Oct'19	Nov'19	Dec'19
1 NGT 1.0 Deployment	11.4	10.4	12.1	13.7	15.2	15.6	17.8	18.2	20.5	20.9	20.9	20.9
2 NGT 1.1 Deployment	10.2	9.1	10.5	11.2	11.2	11.3	13.0	13.4	16.1	17.0	18.5	18.5
3 Cost Reduction Fab4 - Athena	6.5	5.5	5.0	4.5	4.2	4.7	5.9	6.3	9.0	9.9	11.4	11.8
4 NGT 2.0 Development - BE	4.8	3.6	3.1	2.6	2.3	2.4	3.6	4.0	6.7	7.6	9.1	8.7
ZBB line												
4 NGT 2.0 Development - FE	1.3	1.1	-0.4	-0.7	-1.0	-0.9	-0.2	0.2	2.9	3.8	5.3	5.3
5 Cost Reduction Fab3 - Aurora	0.4	-0.3	-2.0	-2.2	-1.2	-0.9	-0.2	0.2	2.9	3.8	5.3	5.3
6 Cost Reduction Fab3 - NGT 1	0.4	-0.3	-2.0	-3.7	-2.8	-2.8	-1.9	-1.5	1.2	0.9	2.3	2.3

Figure B3. SunPower's CTD ZBB worksheet. The cells show remaining engineering capacity after each priority program is added in rank order. The horizontal red line marks the "ZBB Line" — the boundary between programs that fit (above) and programs that do not (below). Cells in pink indicate negative remaining capacity. The line crosses priority 4 — half the program fits, half does not. Priorities 5 and 6 are entirely below the line.

Two observations about the worksheet are worth making. First, the line is not at a single program boundary — it cuts horizontally through priority 4 (NGT 2.0 Development), with the back-end (BE) sub-program fitting and the front-end (FE) sub-program not fitting. This is a common pattern in real organizations: programs decompose into sub-programs that have different priority, different staffing, and different durations, and the practical line falls between sub-programs rather than between programs.

Second, the negative-capacity cells in the lower right of the worksheet (pink-shaded) are the explicit, visual price of attempting to start priorities 4-FE, 5, and 6 in addition to priorities 1-3. The worksheet does not just say "you cannot do it" — it shows by how much. In February it shows the team would need 3.7 more engineers than it has, in April 3.7 again, and these gaps persist through August. This is the visualization mechanism that makes the cost of overcommitment undeniable in the room: no executive can argue that 23 engineers is enough to staff -3.7 more positions.

The mechanism for closing the gap is what SunPower called pipelining: a structured per-engineer review in which lower-priority work is explicitly deferred, descoped, or moved to later quarters until the monthly capacity comes back into balance. The pipelining process took each engineer's assignments and stretched the lower-priority programs across more months at lower per-month allocation, until the cumulative demand stayed under the available-capacity ceiling. The output was a realistic plan that the engineering leadership could commit to.

Key finding. The ZBB Line is the most direct visual expression of starts control available. Programs above the line are funded, programs below are deferred, and the line moves only when capacity changes or priorities change. The pink cells are the price of pretending the line is somewhere else than where the math actually places it. Making the line visible to every leader in the organization is what converts a theoretical capacity argument into an operational decision.

Section B.3

How the SunPower discipline maps to the FTTM model

The mapping from the SunPower implementation to the discipline described in the body of this paper is direct. The AHP-style unconstrained ranking from section three is the priority ordering of the six CTD programs (NGT 1.0 deployment first, NGT 1.1 deployment second, cost reduction Fab4 third, NGT 2.0 development fourth, cost reduction Fab3 programs fifth and sixth). The constraint application from step six of the AHP process is the ZBB worksheet. The starts control gate from section four is the explicit ZBB Line where cumulative demand crosses the 23-engineer capacity ceiling.

The decision-rights model from section five maps onto the cross-organizational SG approval board. The board's nine functional areas are the input and agree roles in the IRADP framework. The general manager, sponsor, or technology lead serves as the named Decider depending on the gate — for SG2 plan-of-record decisions the sponsor is the Decider, for SG4 limited-release decisions the SG approval board votes by super-majority with the CEO or general manager casting the deciding vote in event of a tie.

The change-control mechanism on the plan-of-record at SG2 is the SunPower analog of Altera's emergency-review triggers. Both processes name the conditions under which an active program returns to governance for re-evaluation, both require those conditions to be defined when the program is approved (not retroactively when something breaks), and both keep the bar high enough that the trigger fires rarely but predictably when it does. The two implementations differ in detail — Altera tracked triggers explicitly inside fastDecisionAI, SunPower tracked POR change reviews inside the SG documentation system — but the discipline is the same.

Key finding. The two case studies in this appendix are independent confirmations of the same pattern. Altera and SunPower built their portfolio disciplines in different decades, in different industries, and without reference to each other. Both arrived at the same essential structure: unconstrained AHP-style ranking, explicit constraint modeling against real engineering capacity, named decision-rights with a single accountable Decider, and a change-control discipline that protects commitments without freezing them. The pattern is what works.