069 The Decision Architecture Overview

Perspective Book · https://kb.jeffreysummers.com/docs/069-the-decision-architecture-overview/

Most operators make their most important decisions the same way they make their least important ones -- quickly, from gut feel, with whatever information happened to be available at the moment the decision could no longer be avoided.

That works for the small decisions. The table that needs to be turned. The cast member who needs a redirect. The dish that needs to come off the pass. Those decisions are made in seconds and should be. The read, the judgment, the experience -- all of it fires in real time and produces the right call without a formal process.

The significant decisions are different. Changing the broadline vendor. Re-engineering the menu. Making a key hire. Changing the service model. Restructuring the operation. These decisions touch every domain simultaneously. They produce consequences that compound for years. And most operators make them with the same gut-feel process they use to decide what to run as a daily special.

That is not a character flaw. It is the absence of architecture.

The Decision Architecture is a complete system for making significant decisions -- not faster, but soundly. It does not slow the operator down. It makes sure the resources deployed against a significant decision are pointed at the actual problem rather than the symptom the operator noticed first. It has a beginning, a middle, a gate, and a verdict. Every component connects to the next. You do not get to skip the gate because you are confident. You do not get to skip the verdict because the decision felt right.

The system runs through [The Operator Decision Tree]. Seven layers. Each layer has a job. None can be skipped without corrupting the layers downstream.

Before the Tree runs -- the precondition context

Every decision the operator makes sits inside one of two response patterns.

[Reactive Dangers] is the default. The decision-making process runs only when externally triggered -- by a known decision point the situation requires (opening a restaurant so a menu is

needed, hiring so an interview gets scheduled) or by an unknown decision point that forces itself into view (the number moved, the complaint arrived, the competitor took the ground). Known or unknown, the trigger is always external. If nothing declares itself loudly enough to force a response, nothing gets examined. The gaps that never announce themselves compound unobserved. The opportunities that never package themselves get found by someone else. [Static Decline] runs in the silence between triggers.

[Proactive Read] is the discipline that replaces it. The standing examination of current state against MDV potential -- running continuously, across all five fundamentals and all five stakeholder surfaces, whether or not anything has declared itself. The operator running [Proactive Read] is not waiting for a known decision point or an unknown crisis to force the process. He is always looking for the decision that needs to be made before it announces itself.

The Decision Architecture was built for the operator who has adopted [Proactive Read] as his standing posture. It can be run reactively -- most operators do, most of the time. But it was designed for the operator who runs it on purpose, before the trigger arrives, because he is always looking for what needs to be decided whether or not anything required him to look.

The Seven Layers

Layer Minus 1 -- The Standing Examination

The continuous read of current state against MDV potential. Runs permanently across all five fundamentals and all five stakeholder surfaces. Not a scheduled review. [Proactive Read] in operation. Surfaces gaps before they become crises. Surfaces innovation candidates before they package themselves as obvious options. When a gap is identified: does the operator have sufficient sight to decide, or is evidence required first? Sufficient sight -- the Tree runs from Layer 0. Insufficient evidence -- Exit C is already indicated before Layer 1 runs.

Layer 0 -- Declaration

Precondition. Runs before the Tree begins on any identified gap or decision. The operator names his priors -- what he already believes about this decision. Names his bias position -- confirmation bias, [Golden Rule Bias], or [Disruption Bias]. Declares the evaluation standard -- what a clean verdict looks like before the evidence arrives. [Bias Prosecution] cannot run without a named defendant. Declaration is how the defendant gets named. The operator who skips Declaration is

running [Bias Prosecution] without a prosecution target -- which means the bias runs the verdict unchallenged.

Layer 1 -- The Anchor

MDV at the top. Every decision measured against it. [Operator's Lens] directly beneath -- aggregated perception from every role in the building. Not the operator looking at the room. The operator reading everyone else's read of the room. Not superior perception. Aggregated perception.

Layer 2 -- The Inputs

Ten Sight Targets. Eight internal: Forward, Present, Past, Lateral, Vertical, Time, Self, Cast. Two external: Guest, Orbit. Not every decision needs all ten. Relevance decides which targets a given decision requires. The gaps in the targets are the gaps in the sight. Fill them before the gate.

Layer 3 -- Aggregation

Where sight gets synthesized into a coherent picture ready for the gate. Also where decisions go to die -- five Pipeline Failures live here, and if they run unchecked long enough they collapse into [Vision Swampacolypse]. [By Design Or By Default] runs underneath the entire layer.

Six concurrent stress-test tools run inside Layer 3. Not sequential. Concurrent lenses applied until the picture is complete enough to push to the gate.

The Five Questions screen whether the decision belongs in the operation at all. [Structural Scale] tests whether the operation can actually execute it. [Five Stakeholder Read] evaluates the decision against all five surfaces simultaneously -- business, cast, Guest, community, operator -- to surface trade-offs before they become consequences. No single surface kills the decision alone. The Decoy Option test names the option that exists to make another look better and removes it from the evaluation. [Bias Prosecution] runs the aggregated read against the declared priors from Layer 0 -- where evidence confirms the prior, note it; where it contradicts, that contradiction is the signal; strip the preference, strip the historical scorecard, strip the disruption lean; just the facts. [By Design Or By Default] runs underneath all five as the standing discipline.

[Read Log] runs alongside aggregation as a standing capture discipline. Any notable Layer 3 output worth carrying forward gets logged before the gate runs. The read that is not captured is

the read that never happened.

Layer 4 -- The Gate

[Burn The Boats]. Two questions, fixed order. Does this decision expand the integrity of the MDV? Is this decision a force multiplier for the MDV? Both must return yes.

Three exits -- not two.

Exit A -- both yes: Burn. Committed, irreversible push. No retreat, no half-commit, no second-guessing once the boats are gone. Moves to Layer 5.

Exit B -- either no, sufficient sight: Decision returns to Layer 3 for more aggregation. The gate does not open on one yes.

Exit C -- insufficient evidence to gate: The decision cannot be gated soundly without a controlled test. The operator has sufficient sight to know a decision is required but insufficient evidence to gate it honestly. Pilot required. Triggers Layer 4B -- [Operator R&D]. Does not return to Layer 3 yet. Evidence has to be built before aggregation can complete. Exit C is not a failure of the gate. It is the gate doing its job -- recognizing that burning on insufficient evidence produces a committed decision built on assumption. The pilot builds the evidence the gate needs to burn cleanly.

Layer 4B -- [Operator R&D]

Runs only when Exit C triggers. The formal pilot process. Ten steps. The discipline is in the completion. The operator who reaches Exit C has sufficient sight to know a decision is required and insufficient evidence to gate it. [Operator R&D] builds the evidence. The gate then runs on evidence, not assumption. The full process is taught in its own section. The loop: pilot runs, evidence re-enters Layer 3, aggregation runs again, gate runs, Exit A or Exit B. No second Exit C on the same decision.

Layer 5 -- The Push

Forward decision made and committed. Feeds [Outcomes Formula], Learn-Coach-Relearn Paradigm, and Outcome By Design. The post-action phase of [The Pause Principle] runs the result backward: the thinking behind an outcome is only as good as the outcome it produced. The outcome is the verdict on the thinking. If the outcome is right, the thinking was right. If the outcome is wrong, something in the architecture broke -- and the post-action read names where.

The Full Loop

Standard decision: Declaration → Pause → Sight → Aggregate → Gate → Burn → Push → Verdict → Pause → Examination continues.

Exit C decision: Declaration → Pause → Sight → Aggregate → Gate → Pilot → Re-enter Aggregate → Gate → Burn → Push → Verdict → Pause → Examination continues.

Standing underneath both loops: [Proactive Read] running continuously -- surfacing gaps, generating candidates, routing them into the Tree before [Reactive Dangers] forces the hand.

What Changes Tomorrow

Pull up the Tree the next time you catch yourself circling a significant decision -- not the daily calls, the ones that will still have consequences a year from now. Before anything else, name out loud which layer you are actually in. Name your bias position before the read begins. Run all six Layer 3 tools before the gate. And ask at the gate whether Exit C applies -- whether the decision requires evidence you do not yet have before it can be burned cleanly. Most operators discover they have been stuck in Layer 3, still aggregating, long after they believed they were ready to push. Some discover they have been burning on insufficient evidence when what the decision required was a pilot. Find your layer first. Everything else follows.