



# Speed to Value in Aviation IT

Choosing the Right Delivery Model  
When Programs Can't Slip

---

For Aviation IT & Engineering Directors, Program Owners, and Delivery Leads

SECTION 1

## The Delivery Reality

Aviation IT programs aren't abstract delivery challenges. They're mobile crew applications that have to work before the first flight of the day. They're cloud migrations threading through legacy reservation systems that can't be taken offline. They're API modernization efforts connecting operational data to passenger-facing surfaces in real time, cybersecurity programs running against fixed regulatory calendars, and passenger-facing technology deployed across multiple live airport hubs without disrupting a single check-in.

The programs don't pause. The terminals don't close.  
The delivery dates are rarely the kind that move.

If you're reading this, you've probably already said yes to a date in a room full of people. You own it. The question isn't whether to staff the program; it's whether the delivery model you're considering will protect the timeline you've committed to.



### ONE GOVERNING PROMISE

This guide helps you choose the delivery model that protects velocity when programs can't slip. It's a decision framework to help match the model to the moment before the commitments are locked in.

## SECTION 2

# The Two Clocks That Define Speed to Value

Most organizations measure staffing speed by time-to-fill and time-to-value (or productivity). Both matter, but neither captures where the delays actually concentrate. To evaluate speed the way delivery experiences it, you need two clocks running at the same time.



## Clock 1: Time to Fill

**Need recognized → Candidate surfaced → Selected → Cleared → Starts.**

This is where most hiring metrics live, and most organizations manage it well. It's not where most programs lose time.



## Clock 2: Time to Productive Value

**Start date → Access confirmed → Context established → First useful output delivered.**

This is where most delay occurs. A role filled does not immediately lead to work unblocked, so there's an increased focus on how to speed the worker into increased, productive output.

Each model shifts risk differently across these two clocks. The question is not which model is best in general, but which one best protects the timeline in front of you. The next section breaks down how each model behaves and when it makes sense to use it.

*A model can be right in general and wrong for a moment. The job is to match model to moment before the moment arrives.*

SECTION 3

# How to Choose the Right Model

The three models below are decision paths, not examples. Each is right in specific circumstances and carries real tradeoffs on the two clocks. The point isn't to rank them; it's to make those tradeoffs visible early, before the wrong choice shows up as a delivery problem.

Each model is evaluated the same way: when it protects velocity, how it behaves on both clocks, the risks most likely to cause slip, and what to confirm before committing.

SECTION 3A

## Permanent Hiring

When Ownership and Continuity Matter Most

Permanent hiring is built for the long game. It's right when the capability need extends well beyond any current milestone and the role requires genuine organizational embedding to work well. The risk isn't the model itself; it's choosing it against a fixed deadline when the hiring clock and the delivery clock are running at different speeds. That gap is manageable when it's planned for. It's expensive when it isn't.

### WHEN IT PROTECTS VELOCITY

- Ongoing capability needs with steady-state demand and no acute phase dependency
- Roles tied to long-lived ownership: platform stewardship, core systems, team leadership
- Critical "required to operate" applications where institutional knowledge walking out is an unacceptable risk
- Situations where an established internal pipeline can realistically meet the delivery timeline

### HOW IT BEHAVES ON THE TWO CLOCKS

#### CLOCK 1: TIME TO START

Typically the slowest of the three models. Full-cycle recruiting, leveling and compensation alignment, interview loops, background checks, and notice periods all extend the gap. For specialist aviation IT roles, six to twelve weeks from opening to start is realistic, often longer. The risk is choosing this model when the delivery window doesn't have that time.

#### CLOCK 2: TIME TO PRODUCTIVE VALUE

Variable, and often underestimated. A new permanent hire arrives with strong long-term potential and a real ramp ahead of them: system access, program context, relationships, and orientation to where things currently stand. When onboarding is well-prepared, that ramp is short. When it isn't, even a strong hire spends the first weeks producing less than the delivery plan assumed.

## VELOCITY RISKS TO WATCH

- The search extends toward the ideal candidate while the window narrows. Hiring cadence doesn't compress to match a fixed delivery date.
- Internal approvals for leveling and compensation add time that's invisible in the external recruiting timeline.
- A late-phase arrival creates rework risk downstream; teams that were sequenced behind this role have been waiting.
- When the hiring manager is deep in delivery, interview bandwidth becomes a hidden bottleneck.

### RISK SIGNAL

If the start date is being set by the hiring calendar rather than delivery sequencing, permanent hiring isn't protecting velocity for this moment.

## WHAT TO VERIFY BEFORE COMMITTING

- ✓ Is the capability need persistent, or does it end with the current milestone? Persistent needs justify the lead time. Time-bound ones usually don't.
- ✓ Can recruiting realistically reach the required start date — the one built from delivery sequencing, not the one that sounds reasonable?
- ✓ Are credentialing and background check steps scheduled and owned by a named person before the search begins?
- ✓ Does the hiring manager have real bandwidth to move interviews at delivery pace?

### SECTION 3B

## Staff Augmentation

When Targeted Speed and Flexibility Matter

Staff augmentation is a precision instrument. It adds the right capability at the right moment without changing internal governance, ownership, or direction, and when it's set up well, it's genuinely fast. The catch is that fast to add and fast to impact are not the same thing. Augmentation can move Clock 1 and stall Clock 2, and the difference almost always comes down to how prepared the receiving environment was before the specialist arrived.

## WHEN IT PROTECTS VELOCITY

- Targeted skill gaps (mobile, API, cloud, data, security) where the internal team needs coverage, not direction
- Scenarios where the internal team can carry the program forward after the milestone but needs reinforcement to reach it
- Situations requiring rapid ramp without changing internal ownership or governance
- Programs that need to scale capacity up or down across workstreams without renegotiating the whole structure

## HOW IT BEHAVES ON THE TWO CLOCKS

### CLOCK 1: TIME TO START

Generally faster than permanent hiring when the supplier pipeline is active and procurement is clear. Specialist contractors in aviation IT disciplines can often be sourced and confirmed in two to four weeks. But vendor onboarding, VMS workflow, rate approvals, and contract routing can add that time back invisibly, and interview processes that replicate a full permanent hiring cycle can cancel the speed advantage entirely.

### CLOCK 2: TIME TO PRODUCTIVE VALUE

The highest-risk clock for this model, and almost entirely within the client's control. A specialist who starts on time into an unprepared environment (no documented dependencies, no named onboarding owner, no defined week-one output) isn't productive on day one regardless of how fast the hiring moved. The environment determines Clock 2, not the candidate.

## VELOCITY RISKS TO WATCH

- Onboarding falls to internal leads already running at capacity. Adding people slows the team before it speeds them up.
- When it's not clear what the specialist owns versus what the internal team must drive, both sides wait for the other to move.
- The specialist is ready, but stuck waiting on decisions, access, or context they have no way to unblock themselves.
- VMS or procurement steps that weren't mapped push start dates past the window they were meant to protect.
- Undocumented tribal knowledge means each new specialist starts context-building from zero.

### RISK SIGNAL

If onboarding and access aren't ready before day one, augmentation looks fast at the offer stage and slow at the impact stage. The hiring process worked; the delivery window slipped anyway.

## WHAT TO VERIFY BEFORE COMMITTING

- ✓ Who owns onboarding and access setup (a named person, not "the team")?
- ✓ What does good output look like in week one and week two? Not "getting up to speed;" something specific enough to point to.
- ✓ Are dependencies documented at a level someone new can actually use? Not architectural diagrams, but "here's what you'll touch and who owns each piece."
- ✓ Can the internal team absorb the coordination load this model creates right now, given what else they're carrying?

SECTION 3C

## Outcome-Based Delivery

When Capacity Without Internal Lift Is the Priority

Outcome-based delivery is used when work can be clearly bounded and the program needs capacity without adding internal coordination load. The team owns the scope and is accountable for milestones, not for filling seats. When this model is set up properly, it's often the fastest path to output. When it isn't, it's the model most likely to ramp quickly and then stop. A capable team arriving into unclear scope, undocumented dependencies, or missing access can't deliver its way out of problems the client organization hasn't solved.

### WHEN IT PROTECTS VELOCITY

- Outcomes that can be clearly bounded: a discrete capability, an integration milestone, a release-ready component with defined acceptance criteria
- Programs where internal teams need capacity without absorbing more coordination overhead
- Integration or train-the-trainer scenarios: specialists upskilling an internal team while delivering against the platform
- Greenfield or standalone workstreams where dependency surfaces are well-documented and system complexity is bounded

### HOW IT BEHAVES ON THE TWO CLOCKS

#### CLOCK 1: TIME TO START

Scoping, commercial definition, governance setup, and integration alignment all happen before the team moves. When that front-end work is done well, the team starts fast and self-directs from day one. When it's rushed, the team starts on time but can't move — which tends to get misread as a slow Clock 2 when the real problem is an incomplete Clock 1.

#### CLOCK 2: TIME TO PRODUCTIVE VALUE

When scope and dependencies are clean, outcome-based teams can close Clock 2 faster than augmentation; they don't need internal decisions to operate. The failure mode is undefined acceptance criteria or undiscovered dependencies that stop the team mid-delivery after the early momentum has been spent.

### VELOCITY RISKS TO WATCH

- Undefined scope consumes build time. The team has capacity and no clear place to apply it.
- Dependencies that weren't surfaced before start become mid-delivery blockers with no established path to resolution.
- When acceptance criteria are ambiguous, completion becomes a negotiation rather than a confirmation.
- No named internal owner for integration decisions means the team waits on answers no one has been empowered to give.
- Stakeholder alignment takes longer than planned and compresses build time before the team has properly moved.

**RISK SIGNAL**

If the organization can't define done and unblock dependencies before the team starts, this model will ramp quickly and stall midstream. The team is capable; the setup wasn't complete.

**WHAT TO VERIFY BEFORE COMMITTING**

- ✓ Can the definition of done be stated in plain language right now — in terms the team can act on, not as a diagram?
- ✓ Are the systems and interfaces the team must touch identified, documented, and accessible?
- ✓ Is there a named internal owner for integration decisions, someone available and empowered, not just nominally assigned?
- ✓ Is success defined as outcomes and milestones, not as people or hours on the engagement?
- ✓ Is governance (reporting cadence, escalation paths, acceptance gates) established before the team starts?

SECTION 3 — REFERENCE

# Speed Impact Grid

A quick reference for planning conversations. Where each model is fastest, where time tends to hide, and what to confirm before committing.

MODEL	FASTEST WHEN...	HIDDEN LEAD TIME	FRICTION RISK	VERIFY FIRST
<b>Permanent Hiring</b>	Ongoing capability; long-lived ownership roles; stable pipeline for role type	Leveling & comp alignment; interview loops; background checks; notice periods	Hiring cadence mismatched to fixed windows; late phase arrival	Start date realism; credentialing scheduled; hiring manager bandwidth confirmed
<b>Staff Augmentation</b>	Targeted gaps; rapid ramp with stable internal ownership; elasticity needed	Vendor onboarding; VMS & procurement; rate approvals; access & tooling	Internal lift; ownership ambiguity; dependency waits; coordination overhead	Named onboarding owner; week-one output defined; dependencies documented; VMS lead time mapped
<b>Outcome-Based Delivery</b>	Bounded outcomes; capacity without internal lift; self-starting teams required	Scoping & commercial definition; governance setup; integration alignment; stakeholder approvals	Unclear scope consuming build time; dependency blocks; undefined acceptance criteria	Plain-language definition of done; integration owner named; access confirmed; governance in calendar

## SECTION 4

# Hybrid Reality: Combining Models Without Losing Control

Complex aviation IT programs rarely run a single model start to finish. Mobile, cloud, integration, data, and security don't stay in separate lanes; as integration surfaces constraints, demand shifts in ways that no single access model handles cleanly across every phase. The organizations that protect velocity through complex programs aren't the ones that found the best model; they're the ones that planned the right combination before the pressure arrived.

Hybrid approaches tend to emerge in one of three situations: a model that works well for one phase is structurally too slow for another; hidden lead time surfaces late in one workstream and a parallel approach is needed to protect the overall timeline; or scope shifts in ways that one model can flex to handle and another can't. In all three cases, the answer is the same: anticipate the combination and design the handoffs before you need them.

## HOW TO SEQUENCE MODELS ACROSS PHASES

Work backward from the high-stakes moments. Each moment has a speed profile requirement, and the model should match it. Common patterns in aviation IT programs:

- Outcome-based delivery for a bounded integration milestone in Phase 1, transitioning to staff augmentation for ongoing support as internal ownership stabilizes
- Permanent hiring running in parallel with augmentation: augmentation holds velocity during the critical window while the permanent hire completes the full recruiting cycle
- Outcome-based delivery for a greenfield build, augmentation for the cutover window, permanent hiring for post-launch stewardship

## THREE RULES FOR CLEAN HANDOFFS

Friction in hybrid approaches doesn't come from the combination; it comes from the transitions.

**Define ownership boundaries explicitly.** Which model owns which workstream or phase needs to be stated out loud. The assumption that everyone knows is where overlap friction starts.

**Plan knowledge transfer as a delivery requirement.** Especially when transitioning from outcome-based delivery to internal or augmented teams. What the outgoing model knows and what the incoming model needs are only the same if someone makes them that way.

**Confirm the receiving model's readiness before the handoff, not during it.** Access, tooling, and documentation need to be in working shape before whoever is picking up the work arrives. A handoff into an unprepared environment isn't a transition; it's a restart.

SECTION 5

# Which Model Fits Your Moment

Four questions. Answer based on your constraints right now, not on what's worked before or what feels familiar.

**Q1 — Is there a delivery date, phase milestone, or compliance window that can't move?**

- No →** The timeline has flexibility. Permanent hiring may be the right long-term investment. Evaluate role persistence in Q2.
- Yes →** Time-to-start matters. Move to Q2.

**Q2 — Is this a persistent capability need, and can the timeline absorb full-cycle hiring?**

- Persistent, 8+ wks →** Permanent hiring is likely right. Verify pipeline realism against Section 3A. If a critical phase arrives before the hire is in place, run augmentation as a bridge.
- Persistent, <8 wks →** Permanent hiring makes sense long-term but won't protect this delivery window. Run augmentation now while the permanent search runs in parallel.
- Time-bound →** Permanent hiring is the wrong model for this moment. Move to Q3.

**Q3 — Can the work be clearly bounded, or does it need close integration with the internal team?**

- Bounded →** Outcome-Based Delivery is a strong candidate. Can you define done in plain language, name an integration owner, and confirm access before the team starts? If yes, proceed. If not, resolve those gaps first. See Section 3C.
- Deeply embedded →** Staff Augmentation is the better fit. Verify onboarding readiness, internal lift capacity, and VMS lead time before committing. See Section 3B.

**Q4 — Are you running multiple workstreams with different characteristics?**

- Yes →** Different workstreams will likely need different models. Run Q1 through Q3 per workstream, then use Section 4 for sequencing and handoff guidance.

SECTION 6

# Readiness: The Second Half of Protecting Velocity

Choosing the right model is the first half. Being ready to execute it is the second, and it's where most of the recoverable time lives.

The verify lists in Section 3 are model-specific. What follows applies to all three and determines whether Clock 2 closes quickly or slowly once the engagement begins.

## READINESS PACK — CONFIRM THESE BEFORE DAY ONE

### **Access path defined and owned**

Systems, repositories, tools, and environments confirmed with a named owner before day one. Security provisioning and data access steps that aren't scheduled before the engagement starts will compress the first sprint.

### **Dependencies documented at a usable level**

APIs, data sources, and integration surfaces documented at a level the incoming person can act on; not diagrams, but "here's what you'll touch, here's who owns each piece, here's what's known and unknown."

### **First useful output defined**

What good looks like in week one, in specific terms the incoming person can act on from day one. Not "get up to speed;" something that can be pointed to.

### **Decision owners named**

Individuals, not committees. Someone who can answer a question on a Tuesday afternoon without scheduling a meeting. The absence of a named owner is the most common source of first-week stall.

### **Delivery cadence in the calendar**

A standing mechanism for unblocking issues and getting feedback, already scheduled before the engagement starts. Not something to set up after the first blocker appears.

### **Compliance and credentialing scheduled**

Security clearances, background checks, and aviation-specific credentialing steps tracked and owned. These don't happen in the background, and in aviation environments, they surface as surprises at exactly the moments you can least afford them.

*The goal isn't a perfect environment. It's a prepared one, where the incoming person or team can move on day one rather than spend the first sprint discovering what should have been confirmed the week before.*

SECTION 7

# Your Next Tool: Model-Specific Readiness Checklists

This guide helps you choose the right model; the checklists help you execute it. They're different jobs, and the distinction matters.

Each checklist surfaces the lead-time exposures, procurement steps, and friction risks specific to that model, the ones most likely to cause slip between decision and delivery. They're structured for planning conversations *before* commitments are made, not retrospectives after delays appear.

**CHECKLIST A**

## Permanent Hiring Readiness

Time-to-start realities, approval and compensation loops, credentialing timelines, onboarding ownership, and ramp-to-value readiness. Use this to confirm the hiring clock can meet the delivery clock.

**CHECKLIST B**

## Staff Augmentation Readiness

VMS and procurement lead time, access provisioning, internal-lift assessment, onboarding ownership, dependency documentation, and first-week output definition. Use this to find the friction that makes augmentation feel slower than it should.

**CHECKLIST C**

## Outcome-Based Delivery Readiness

Definition of done, governance and escalation path set up, dependency clarity, integration ownership, and access readiness. Use this to confirm the environment is ready for the team to move from day one.

**Download the checklist for the model you've chosen**

Structured for use in planning meetings before decisions lock in.

[lorienglobal.com](https://lorienglobal.com)

SECTION 8

# Protecting Velocity Is a Leadership Choice

The programs that hold their timelines aren't the ones that found the best staffing model. They're the ones where leaders matched model to moment, confirmed readiness before commitments locked in, and designed hybrid structures that accounted for how the program would actually evolve.

Those decisions protect timelines. They protect delivery confidence. And they protect the capital that mismatched choices consume quietly, without ever appearing on a risk register.

That's not a staffing discipline; it's a delivery discipline. And for aviation IT programs that can't slip, it's the one that matters most.



**READY TO TALK THROUGH A SPECIFIC STAFFING DECISION?**

Speak with a Lorien Aviation Delivery Specialist about your current program: the model, the moment, and what readiness looks like for your specific constraints.

[lorienglobal.com](http://lorienglobal.com)