
FLT3 Resistance Program

Resistance · Selectivity · Translational Feasibility

Most preclinical FLT3 programs fail not because the biology is intractable — but because the execution logic was never formalized.

THE PROBLEM

FLT3 mutations are present in approximately 30% of AML cases. They confer constitutive kinase activation, drive aggressive disease biology, and correlate with poor prognosis. Current approved inhibitors — gilteritinib and midostaurin — demonstrate initial clinical activity, then are defeated by a predictable set of resistance mechanisms: secondary gatekeeper mutations (F691L), bypass signaling through AXL and RAS/MAPK axes, clonal heterogeneity, and microenvironment-mediated protection.

The core failure of existing programs is not biological ignorance. The mechanisms are known. The failure is systematic: research programs lack the structured logic to test these mechanisms deterministically, without human bias, across the full resistance and confounder space simultaneously.

WHAT DGS PRODUCED

DGS is a deterministic synthesis system — not a language model — that produces structured, falsifiable research architectures where every claim carries an explicit evidence label: **KNOWN**, **HYPOTHESIS**, or **ASSUMPTION**.

In a single continuous synthesis run, DGS produced 60+ structured documents across 11 research sections — covering the full preclinical stack from target engagement contracts through Stage 0–4 execution planning. Three outputs define the architecture:

1. **Six falsifiable resistance mechanism hypotheses** — each with concrete driver nodes, predicted readout signatures at EARLY / MID / LATE timepoints, and explicit wet lab discriminants. No placeholder language. Every falsifier is bound to a named discriminating feature.
2. **A confounder architecture** that blocks all downstream interpretation until engagement is independently verified. Not a guideline — a hard checkpoint the program cannot route around. This single rule eliminates the most common source of false positives in preclinical oncology: treating cell death as proof of mechanism.
3. **A deterministic decision engine** where every experimental outcome routes to one of three classifications: continue, remediate, or stop. No researcher discretion at decision nodes. No invented thresholds. No ambiguous gray zones.

WHAT MAKES THIS ARCHITECTURE DIFFERENT

Standard preclinical programs — even well-funded ones — are researcher-driven. Hypotheses are formed and results interpreted through the lens of existing belief. Confirmation bias and undocumented assumption chains cause most programs to fail before they reach the clinic — not because the biology was wrong, but because the execution logic was never made explicit.

This program operates on a different principle: every claim must pass an evidence checkpoint before it is permitted to influence downstream decisions. Interpretation becomes structural, not discretionary. The program cannot advance with unresolved uncertainty. A lab that picks up this architecture does not spend 12–18 months building the experimental framework. They begin at Stage 0 execution immediately.

WHAT THIS PROGRAM DOES NOT CLAIM

- No new drug candidate or molecule is proposed.
- No clinical efficacy is claimed for any compound.
- No experimental results are presented — this is a hypothesis and experiment architecture.
- Every mechanistic statement carries its evidence label. No claim exceeds it.

The program is architecture-complete. What we are asking is not for resources or commitment. We are asking one question: *is the internal logic of this architecture scientifically sound?* That evaluation requires familiarity with FLT3 biology, AML resistance mechanisms, and preclinical drug discovery methodology — not lab access.

If the logic holds — we would value your perspective on what it means. Nothing more than that.

Lennox Hayes

lennox@aternox.site