

FIG. 4 – DATA FLOW: SOURCES → SYNC WRAPPER → FIRESTORE → DASHBOARD → OUTREACH

Risk as a weighted five-factor score — not a single metric

Context. A single number produced false positives (slow but solid students flagged red) and false negatives (fast but guessing students flagged green).

Decision. Risk Score = Debt Exposure 30% · Velocity 25% · Precision Decay 20% · Stability 15% · Stall 10%.
Thresholds: ≥60 **RED** · ≥35 **YELLOW** · <35 **GREEN**.

Accepted trade-off. Harder to explain in one line, but far more faithful to the student's reality — and fidelity is what makes a mentor trust the color.

07

BUILD & VALIDATION

AI-FIRST PROCESS, PUT TO REAL DATA

Development ran with AI as the primary co-developer —architecture, implementation, and refactors—, preceded by a proprietary notebook that explored the API before building the production wrapper. The pattern was constant: **the decision of what to measure and why was always my own**; AI accelerated the *how to build it*. Built in three and a half weeks of sustained work, with thresholds centralized in a single configuration file to recalibrate without touching the logic.

The dashboard was connected to the full population —~1,600 students— and emitted **automated daily digests** of tier changes (one day it reported 79 changes; another, 43), confirming the live pipeline. It was presented to

the data platform team in a *data insights* session, and the calibration of the five factors was tuned against the cohort's observed behavior.

08

THE RESULT STATUS AND LEGACY

In production, connected to the real population, with automated daily digests. The product didn't stay a demo: it became the visibility layer on which the next product in the portfolio was designed – the Student Performance Tracker [Ch. 02](#), which adds pattern detection and automatic intervention.

Students in coverage	~1,600
Build	AI-first · <i>what to measure</i> = my own decision · sustained pace (~3.5 wk)
Integrations	Practice API · Firebase · Slack · Vercel
Views	Triage · Matrix · Heatmap · Log · Student Modal
Status	In production


LEARNINGS

- **Technical.** The bottleneck of a live dashboard isn't the interface — it's the robust wrapper over the third-party API: incremental sync, retries, fault tolerance.
- **Product.** A well-weighted composite metric beats five loose ones. The user doesn't want five numbers; they want to know who to attend to first.
- **Pedagogical.** Operationalizing learning science into metrics makes actionable what used to be intuition — and leaves a traceable decision trail.

Visibility where there was none

MATH ANALYTICS COMMAND CENTER

The academic status of ~1,600 students in a single view – from manual, one-by-one inspection to an intervention decision in under thirty seconds.



STATUS	STACK	DATA	METRICS
Production	Next.js · Firebase	Practice API	5 risk axes

01

THE PROBLEM

A SYSTEM WITHOUT VISIBILITY

Student performance lived scattered across three platforms and could only be checked one student at a time. Knowing how a student was doing meant opening their profile, navigating to *settings* → *documentation* → *progress report* and reading it by hand. Multiplied by ~1,600 students, proactive monitoring was simply unworkable.

No mentor could answer "*how many of my students are at risk today?*" without going through the platform one by one.

There was no way to catch in time a student who was stalled, guessing, or declining in accuracy within the session.

Interventions were reactive — they arrived when a guide reported the problem, not before.

There was no common language: "doing well" or "doing poorly" were qualitative judgments, never measurable or comparable. Every week of undetected stalling was lost progress that wasn't recovered within the session's window.

02

THE OBJECTIVE

AN EVALUABLE STATEMENT

PRIMARY OBJECTIVE

Give every mentor real-time visibility into the academic risk of their students in under 30 seconds, replacing manual one-by-one inspection with a consolidated, actionable view.

As secondary objectives: turn raw API signals into interpretable pedagogical metrics (velocity, knowledge debt, accuracy decay, stability); prioritize automatically — have the system say *who to attend to first*, not just who exists; enable actionable outreach from the same tool; and operate as proactive monitoring without adding manual reporting load to the mentor.

Success was defined up front: answer "who is at risk today?" in under 30 seconds, with 100% coverage of students in a single view and latency resolved by automatic periodic sync.

03

THE ARTIFACT

THE TRIAGE VIEW

The dominant use is a **morning routine**: the mentor opens Triage at the start of the day, reads the red/yellow/green zones, and decides who to contact before the first session. The Risk Score orders the entire cohort by urgency; the color resolves the decision at a glance.

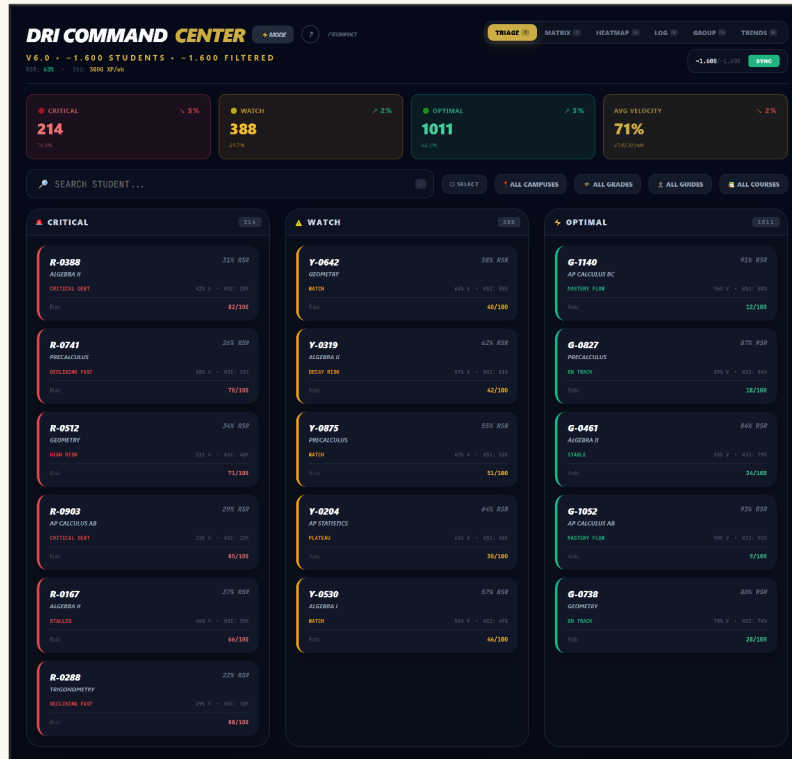


FIG. 1 – TRIAGE VIEW • ~1,600 STUDENTS SPLIT ACROSS THREE ZONES –CRITICAL / ATTENTION / OPTIMAL– AND ORDERED BY RISK SCORE WITHIN EACH ONE • THE COLOR DECIDES WHO TO ATTEND TO FIRST

04

FOUR VIEWS, FOUR QUESTIONS

THE SAME DATA, READ FOUR WAYS

Triage answers "who do I attend to now?". The other views answer different questions about the same cohort: **Matrix** groups by risk pattern —what *types* of student do I have?—; **Heatmap** shows the most critical topics at the cohort level —what *content* is everyone struggling with?—.

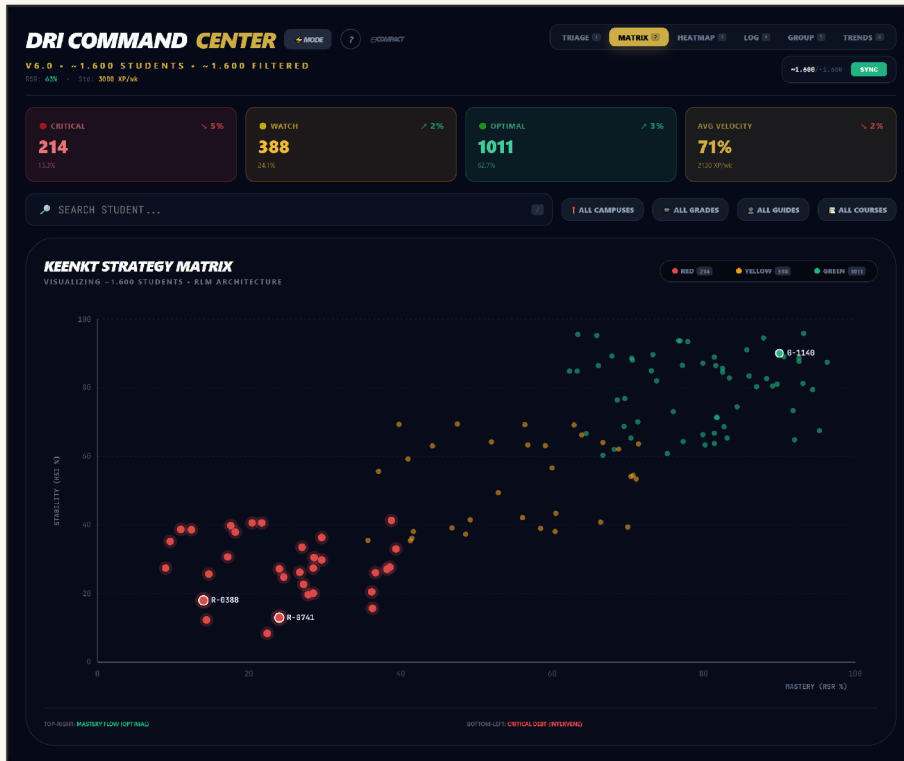


FIG. 2 – MATRIX VIEW • THE COHORT POSITIONED IN THE MASTERY (RSR) × STABILITY (KSI) SPACE, WITH THE UPPER-RIGHT QUADRANT AS MASTERY FLOW AND THE LOWER-LEFT AS CRITICAL DEBT

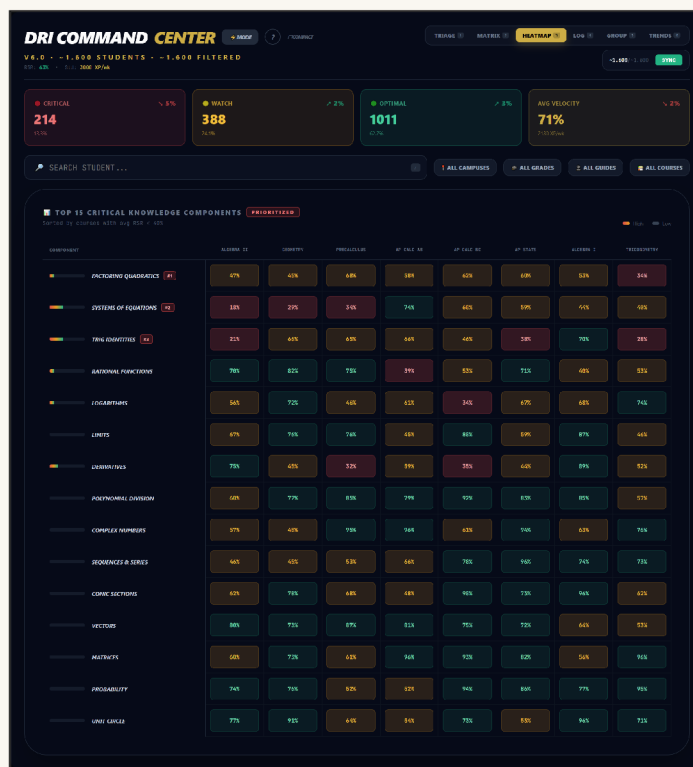


FIG. 3 – HEATMAP VIEW • THE 15 MOST CRITICAL KNOWLEDGE COMPONENTS CROSSED WITH EACH COURSE, PRIORITIZED BY LOW AVERAGE RSR

05

THE FOUNDATION

LEARNING SCIENCE, TURNED INTO METRICS

The dashboard doesn't invent arbitrary indicators: it operationalizes learning-science principles into five measurable metrics, derived from a proprietary calculation protocol documented separately.

PRINCIPLE	METRIC	WHAT IT CAPTURES
Mastery learning	Velocity	The student's pace against the standard of 125 XP per week.
Knowledge debt	DER	Percentage of K-8 topics an HS student is re-mastering — foundational gaps they carry.
Cognitive fatigue	PDI	Accuracy drop between the start and the end of the session.
Retention / stability	KSI	Stability of mastered knowledge over time.
Recent success	RSR	Sustained accuracy across the last ten tasks.

If mentors see risk in real time, they intervene preventively rather than reactively — and the student spends less time stalled within the session's window.

So that no single metric would oversimplify, risk is computed as a **weighted score across five dimensions**, not as a single number. And so as not to distort motivation, the metrics are the mentor's instruments — never a public ranking of the student.

06

THE DESIGN

ARCHITECTURE AND THE DECISION THAT MATTERED MOST

A proprietary wrapper over the practice platform's API handles fault-tolerant incremental sync and feeds Firestore; the dashboard reads from there live, and the message generator pushes outreach to Slack. Stack: Next.js on TypeScript, Firebase/Firestore, Vercel.