

Powering financial aid with ontology

How University of Bridgeport and Goodwin University partnered with CollegeVine to unify Slate and Colleague into one data layer, then rebuilt aid leveraging around the full student lifecycle



THE STARTING POINT

Five connected institutions, solved one project at a time

The University of Bridgeport, Goodwin University, and three magnet schools run on shared resources. Students move between them constantly: a magnet-school pipeline into the universities, approved joint degrees, credits earned at one institution and counted at another.

Bridgeport was not new to AI either. Chatbots were running and voice agents were live. But retention, student success, and every AI project alike were being solved one at a time, and the team didn't have the IT budget to continue doing that.

“I went in thinking they were going to sell me another chatbot. I walked out understanding how I could deploy across the enterprise.”



Sandeep Mannava

Chief Information Officer, University of Bridgeport and Goodwin University

THE PROBLEM

Financial aid is too deterministic

The approach has not changed for the past two decades: set a freshman revenue goal, work backwards into bands by GPA and test score, and tweak the merit award.

The landscape underneath it moved. International enrollment is thinning, high-revenue students are fewer, and more of the class is now non-traditional. The class hits its headcount and revenue can still land two or three million dollars short.


Year 1

is all the traditional model prices on

95%

of Goodwin students are part time

SAME INPUTS, SAME DOLLARS, EVERY TIME



THE INFRASTRUCTURE · SLATE AND COLLEAGUE

Students across different systems are merged as one

Recruitment lives in Slate. Everything after enrollment lives in Colleague, which is also the ERP, finance, and HR system. The same student existed twice, once on each side, with no link between the two.

Partnering with CollegeVine, they have been resolved to one person. A single record now carries what a student was recruited on and what they actually cost and pay, properly informing aid decision needs.

3 days

to move Colleague and
Slate onto the ontology

193

Slate tables pulled in

54

Colleague Ethos
Endpoints


ONE STUDENT RECORD ACROSS SYSTEMS



Every record about a student now points to one place: the actual person. This includes the application and the decision, the academic record, the FAFSA result, each line of aid and the fund behind it, and the cost of attendance. After this project, no subsequent one will require a new data pull.

THE ONTOLOGY

One unified data layer



Modeling the whole lifecycle for a more accurate read on where each aid dollar goes

A conventional leveraging model predicts enrollment from pre-enrollment signals and optimizes first-year net tuition revenue. For a population that is largely commuter, first-generation, and part time, that is the wrong horizon.

Bridgeport and CollegeVine modeled matriculation and retention at the program level across four years, and the cost of serving each student alongside the revenue.


4 years

of revenue and cost per student, per program

\$ millions

in additional lifetime revenue from the same aid budget

WHERE A DOLLAR OF AID ACTUALLY MOVES A STUDENT




THE DELIVERABLE

Leadership can now run the scenarios themselves

Bridgeport got an app on the platform: tuition, class-size cap, athlete floor, and cost-per-student dials on one side, expected enrollees, net lifetime value, and per-student awards on the other.

A CFO came in with a hospital partnership asking for a 25% tuition discount, and no way to say what that would cost. Adding a crew team raises the same question about the athlete roster. Both used to mean weeks of gathering data for someone outside the institution, but now they are settings on a screen.

AID PRICING MODEL



IN PRODUCTION

The recommendation writes itself back into Slate

01

Daily admit list

Roughly 20 admits a day,
scored on lifetime value

02

Award posts to Slate

The institutional award lands
on the record, not in a report

03

Colleague packages it

Existing auto-packaging runs
untouched

04

Counselor in the loop

A window to go higher when
the conversation warrants it

Most aid modeling stops at advice. Bridgeport's runs daily: the model scores the new admits, the institutional award posts back into Slate, Colleague picks it up, and the existing auto-packaging finishes the job. A counselor can still move the number in conversation.

WHAT'S NEXT

Downstream of the offer, on the same layer

Persistence and retention agents

Built on the Colleague data already in the ontology, with LMS data to follow

Aid revisited mid-degree

A second-year adjustment for a student the model flags as at risk

Program-level P&L

Four months of outside work, now scoped at about a week on GL data

Software retired to fund it

Three to four retention products sunset, half the savings returned to the CFO

The lifecycle model already predicts who is at risk. The next set of work acts on it: clearing holds faster, cleaning up degree pathways, and revisiting aid in the sophomore and junior years for students worth keeping.

“I have worked with a lot of technology companies over 25 years. Everybody says six months, a year. These folks talk in weeks.”



Sandeep Mannava
Chief Information Officer, University of Bridgeport
and Goodwin University