Higher Ed Geek Podcast – Episode 307: Will Ballard

Katy Oliveira: Welcome to *Next Practices*—a podcast dedicated to the innovative strategies that drive student success. Hear from transformational higher ed leaders on how they're tackling today's most pressing challenges to make a difference for their students and institutions. I'm your host, Katy Oliveira.

Today, we're bringing you a special guest episode in partnership with the *Higher Ed Geek Podcast*. Our CEO at Civitas Learning, Will Ballard, recently sat down for an in-depth conversation about the real-world application of AI in higher education.

Will has been building AI products for 30 years, and in this episode, he addresses common misconceptions that can stall institutional progress. He shares why you don't need "perfect" data to start making an impact, how assistive and agentic AI can streamline the heavy lifting for advising teams, and how predictive analytics can help institutions intervene before a student even knows they're at risk.

It's a compelling look at how we at Civitas Learning view the partnership between technology and the essential human element of education.

Let's jump into the conversation between Will Ballard and Dustin Ramsdell.

Dustin Ramsdell: Welcome to the *Higher Ed Geek Podcast*, where we explore the impact of EdTech on the student experience with engaging, fun, and relevant conversations that honor the wide range of work happening across the higher ed ecosystem. I'm your host, Dustin Ramsdell. I've been working in and around higher ed for the past decade, focused on student support and digital experience.

Join me every week for discussions with some of the best minds in education technology, bringing perspectives on how digital transformation is reshaping the way we recruit, engage, and graduate learners. You'll hear firsthand accounts of how institutional leaders are developing their strategies, along with stories about how these EdTech tools are being built to foster better student outcomes.

With that, let's get to today's episode.

Dustin Ramsdell: Hello and welcome to episode 307 of the *Higher Ed Geek Podcast*, featuring Will Ballard from Civitas Learning. I really appreciate Civitas's approach to student success and retention—and how they've been integrating AI for a long time.

Will offers a helpful explainer and background on how this work comes together. I also appreciated what he had to say about data quality—specifically, pushing back on the idea that institutions need perfectly clean data before they can start integrating systems and synthesizing information in ways that empower staff and faculty to support students proactively.

Many thanks to Will for taking the time for this conversation, and I hope you enjoy episode 307 with Will Ballard.

Dustin Ramsdell: All right—Will, I'm excited to have you on the show. We'll start as we always do: can you briefly introduce yourself and the work that Civitas Learning does?

Will Ballard: I'm Will Ballard, CEO of Civitas Learning. Civitas Learning is a higher ed SaaS company with a pretty unique value proposition. We take data from the systems schools already have and use it to understand the student outcomes that matter most.

What does that mean in practice? We can predict on-time graduation, grades, academic engagement, and when students may be at risk of dropping out. And it's not that we're interested in prediction for prediction's sake—we're predicting so it *doesn't* come true. If we think someone may drop out, we don't want them to. We want you to be warned early enough to intervene and help that student—whether that's connecting them to tutoring, housing support, financial aid, or other resources that help students stay in school. Ultimately, we help institutions support students all the way through to a degree. It's pretty cool stuff.

Dustin Ramsdell: This is something I'm always passionate about. A lot of my background has been focused on how you fuse strong technology platforms with the day-to-day work of student success coaches, advisors, and others—so they can be more proactive and more informed about what's happening with students.

And especially now, with enrollment challenges and recruiting pressures, I think it's only reinforcing how important it is for institutions to retain and support the students they already have.

At the same time, we're at a moment where institutions have a lot of tools, platforms, and data. The question is: how are you using it? It feels like we've hit a watershed moment—more leaders are thinking about how to harness the data they already have and make these systems work better together.

You've been at Civitas for a handful of years now. Have you noticed this kind of thinking becoming more common among institutional leaders? Any reflections on that shift?

Will Ballard: I think that's right. We engage with presidents, provosts, VPs of enrollment, and VPs of student success and student affairs. We're getting attention at the executive level in a way that's different from years past—when institutional research teams were often the ones most focused on data and analysis, and they were trying to convince everyone else to be more data-oriented.

Now, being data-oriented feels like the default setting, and the question becomes: what do you do with it?

One of the biggest shifts post-COVID is that it drove much broader use of learning management systems. More faculty engaged with the LMS, and institutions got more value from those platforms. The technology itself is useful—but what matters most is the data it generates.

Historically, the data systems schools relied on were things like the student information system—where you'd post midterm grades and final grades. And that might be it. You know a decent amount about the student—what they're enrolled in, their transcript—but the data inputs are very low-frequency, maybe once or twice a term.

Then people say, "If they're getting bad grades, maybe they're going to drop out." But the grade is the measurement—it's not the input, it's the output. If you want to help students, you need to look upstream of the grade—what goes into making that happen.

With the LMS being used more consistently, you can see who's logging in, who's completing homework, who's taking quizzes. And for almost every school we work with, it's surprisingly predictive. You can even see, at the student level and course-section level, who's engaging with other students—through chats and discussion tools inside the LMS.

That means you can know day by day when there's an inflection point in someone's risk—as opposed to finding out after they fail a midterm. Because the reality is, once students are in the "fail the midterm" zone, there's not a lot of coming back. You need to know much earlier to be able to help.

That's probably the biggest difference in how data is coming together now: it's more actionable because it's higher-frequency.

Dustin Ramsdell: Yeah, that's a great point. One big factor is the consistency of how these platforms are being used, and how that creates more informative inputs for predictive analytics. Years ago, on the front lines as a student success coach, you mostly had lagging indicators.

Sometimes students have that visceral realization around midterm time that they're not doing as well as they thought. And sometimes you get that "official" confirmation from grades—it's not just anecdotal check-ins where a student says, "I'm fine."

But the limitations of that approach are clear. Now, with more ingredients coming together—like login behavior, time spent, assignment activity, knowledge checks—you can blend signals and be far more proactive than reactive.

Another big change in recent years is the rapid proliferation of AI across higher ed tech. Can you talk about how you and your team are incorporating AI into your student success platform?

Will Ballard: We're kind of "old school" AI. I've personally been building AI into products for 30 years. The first system I built that was legitimately AI—just a different kind—was in the mid-90s. I built a scheduling system for San Bernardino Unified, a K–12 district. You take all the students, schools, classes, and requests and build a schedule for the entire district. It was a big A* constraint-based local search problem, which is a specific type of AI.

So I've been doing this one way or another for a long time—whether it was called AI, not called AI, or called machine learning. The names keep changing. "AI" has come back around as the label for a while now.

Today, I think there are three helpful categories of AI to think about. Everyone knows about LLMs—that's great. But here are three practical ways to apply AI:

1. **Predictive AI**: taking a bunch of data and making a forward prediction about what's likely to happen, so you can either know what's coming or intervene to change the outcome. It's essentially an early warning system.

2. **Assistive AI**: generating and processing content to support work. Five years ago, most people didn't realize they were using AI—but post-ChatGPT, Claude, Copilot, Gemini, it's become widely accessible and easy to use.

3. **Agentic AI**: this newer category where you can describe what you want done in plain language—English, Spanish, really any language now—and the AI can drive other software to get work done. In a sense, it lets everyone be a programmer without needing to know how to code.

So in the Civitas context, what does that mean?

For **predictive AI**, we've been doing this for almost 15 years. We pull in data from your systems every day, clean it, normalize it—there's a lot of real "data janitor" work involved—and we use that to build predictive models around persistence, on-time graduation, engagement with coursework, academic performance, and even career outcomes. Really, the entire student lifecycle once a student commits to a school. The point is: you can make predictions about where students are headed and then intervene.

For **assistive AI**, we have some "what you'd expect" features inside our advising workflow tools. For example, there's a capability we call **Advisor Notes**.

Almost every advising team has standards for what notes should look like. But the reality is, there's always variability in how consistently those standards get followed because it takes extra effort. And taking detailed notes while talking with a student means you're not fully engaged—you're half typing, half talking.

On Zoom, people got used to transcription: "Oh, that's useful." So we made an **in-person** version. During an in-person meeting—about half of advising meetings—you can click a button on your laptop to record. It transcribes the conversation and generates a summarized note in the format you and your team specify. That lets advisors be present with the student while still getting what they need in the case file. And of course, if you have a Zoom recording, you can do the same thing there too.

I like that because it's technology that lets you be *more* present and more human, instead of acting as an intermediary for the technology.

Then on the **agentic** side, the idea is: "I want software to go out and do a bunch of work for me."

We built a feature that's essentially an AI agent, but the interface isn't a chat—it's a button called **Meeting Prep**. You click it, and it automatically gathers everything you'd typically pull manually before meeting with a student.

A typical advisor will tell you that for a 30-minute session, they might spend as much as 30 minutes preparing—collecting information across multiple systems. So we do that automatically: SIS records, LMS records, everything we know about the student. We bring it together, run it through our prompts (and customers can customize this), and produce a summary of what's going on with the student—and, importantly, what's changed since the last time you met.

Otherwise, advisors are logging into Banner and Canvas and CourseLeaf and digging through five or six systems to prepare. We bring that into one place so you can be up to date quickly.

And when advisors carry caseloads of 300 students, it's not reasonable to expect them to remember every detail about everyone. So the system becomes the memory, and the advisor can focus on the human part.

So those are the three categories of AI and how we apply them—and we're constantly adding more.

Dustin Ramsdell: That's awesome. I appreciate you walking through those examples and giving a clear framework. There's been a lot of blurring around what "AI" even means, so that breakdown is helpful.

And I remember those pain points. In a lot of cases, advisors either go without the prep because it's too time-consuming—so the meeting becomes transactional, like "Tell me what you

need”—or they spend so much time preparing that it limits how many students they can actually meet with, and limits their ability to do anything else.

So these are compelling examples of practical ways to implement AI in student success. And it’s exciting because we’re at a moment where a lot of the conditions that make this possible have finally converged.

Can you share a bit about Civitas’s approach to data system integration—and why it’s critical to providing high-quality support to students?

Will Ballard: We get all the data.

What we’ve found works best is getting access to as raw data as possible from a customer. Most of the time, that means a direct, read-only connection to their databases, then we bring the data in, process it, and clean it.

One conversation I have all the time is: “Our data is dirty. We need to clean it before we can do anything.”

That’s just not true.

I get why people think it. Sometimes data is messy. But the real issue is usually that the data is **hard to use**, not unusable.

We’ve invested heavily in data processing—SQL, automated pipelines, and AI in that pipeline—to validate and normalize data as it comes in. We run hundreds of checks—some AI-driven, some rule-based—to detect errors. We handle cleanup as part of the process.

And honestly, there’s no such thing as perfectly clean data in the real world.

A comparison I sometimes use: people look at LLMs and forget they were trained on the entire internet. Anyone who thinks the internet is clean data… it isn’t. But AI is extremely good at dealing with the messy stream of real-world inputs.

If I could convince people of anything, it would be to stop saying, “Our data is too dirty to start.” It’s not.

Dustin Ramsdell: I want to pause on that, because I think it’s important. People can feel intimidated by this work because they assume they need to do a huge data cleanup effort first.

I like your framing. I even think of it as “lived in.” Like, “My house is dirty”—no, it’s lived in. It’s not a hotel. That’s a powerful point: institutions can start where they are and still move forward.

Will Ballard: A big part of that is because we're using AI and statistics.

If you're doing this in a classic rules-based way—"I'm worried about first-time-in-college students who got a D on the midterm"—then I understand why you'd worry about data quality. Rules are rigid.

But with AI and statistics, the data doesn't have to be perfect. And here's something that sounds odd: I don't need to be "right" in the way people assume.

I don't want to predict that someone will fail—and then they fail. I want to predict that someone is at risk—*and then you do something about it*.

It's not a tragedy if I flag someone as at risk and you meet with them and find out everything's okay. The goal is to identify risk early and create opportunities for human intervention.

There's a human in the loop. We're using technology to better inform people, connect people, and help them solve problems with humans. That's why the data doesn't need to be perfect to be useful.

Another piece is **frequency**. Plenty of schools have good predictive work on the IR side, but they run it in batches—once a term, once a semester.

We run scoring, risk, models, and data feeds every day. Why? Because you need frequency to act.

Not being rules-based helps with messy data. Running daily helps you respond before the end of the term—before you're just admiring the problem after it's already happened.

And I'll add one more: we don't use a national model. We don't plug your students into a generalized stereotype. For each institution, we build models based on your history and outcomes. That means we're not generalizing across regions or institutions that aren't the same.

If your data is messy, that's okay—your outcomes are still there: grades, graduation, program completion. We're looking at your real patterns, not comparing you to a national template.

Dustin Ramsdell: That's great—especially the point about institutions being different. We talk a lot on the show about culture and change management and how even good intentions can hit friction.

But circling back: with all this integration and daily updating, why does the technical work matter when it comes to the advisor actually sitting with a student and trying to provide high-quality support?

Will Ballard: For sure. Having all the data enables things like meeting prep, and we also have a unified student profile. A typical institution might have 200 software systems. Advisors won't log into 200, but they might log into 6 to 20. Bringing the data together on one screen saves a huge amount of work and gives people visibility.

And since you mentioned change management and human factors: yes, we're a software company—but we provide more than software as part of the subscription.

We have Customer Development Directors. Instead of traditional account managers who are friendly but may not know your world, we hired people who've led advising teams, worked in advising teams, and done IR work. This is not paid consulting—it's just part of what we do.

We provide regular trainings and webinars, on-site and virtual training, and a training library. We connect customers to other customers. We host workshops where customers share problems, solutions, and case studies.

We also have a support team because data feeds break. That's just reality. Someone adds a field to Banner and doesn't tell anyone, and something stops working. We fix it.

So it's a managed service. We engage regularly.

And for executives, we have something called **Initiative Analysis**. Think about any program you're running—an English tutoring lab, a free lunch initiative, whatever. You can take a simple spreadsheet—student ID and date—upload it, and the system uses AI to match participants with similar non-participants. It's like creating a synthetic randomized trial without having to do an RCT.

Then you can evaluate outcomes: grades, persistence, on-time graduation, career outcomes. That helps leaders answer: "We're spending money on programs—how do we know what's working?" And even: "Does this outperform other options, like lowering tuition?"

Affordability matters to students, and institutions need visibility into what's effective. We can run that analysis, customers can run it themselves, and we use it during executive check-ins.

So no—we're not just a database or a tool where you type in notes. It's more than that.

Dustin Ramsdell: Yeah. The more time and resources institutions invest in student success, the more important it is to have a full, up-to-date picture of students—so teams can be proactive before challenges become insurmountable.

And it's powerful to be able to see outcomes and improvements—so the teams doing the day-to-day work can validate and celebrate what's working, and institutions can align around the same data rather than just anecdote.

As we wrap up, what advice or call to action would you leave listeners with?

Will Ballard: The biggest thing I'd tell people is: this really can help.

This isn't AI and robots coming for your job. This is AI as a tool that can dive into systems, gather data, and help you do the human engagement part of your job better.

People have tried chatbots—sometimes they work, sometimes they don't. We've done studies where "maybe they don't" ends up being the outcome. Our focus is helping people use data to support human connection: advisor-to-student, student-to-student, student-to-faculty.

If people only wanted content, they could watch YouTube or read a book. Going to school is more than content and information. It's human interaction and connection—and we're focused on features that support that.

And the second point I'll return to is the data question: no, you don't need clean data to get started. Whether you do it with us or someone else, don't put off using data to help students because you think you need to clean everything up first. It isn't necessary. AI can work around imperfect data in a big way.

Dustin Ramsdell: On that note—because I think it's profound and really important—thank you to you and your team for the work you're doing. This is near and dear to my heart, and I'm excited to see what you all continue building and how institutions continue embracing this approach, because it's smart, necessary, and important for this moment.

We'll include ways to connect with Will and Civitas Learning in the episode description. Will, thanks again for joining me today.