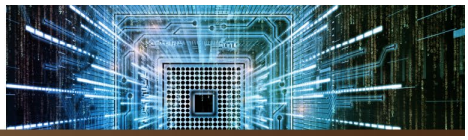

Bentley Systems Inc

Bentley Systems Inc presentation delivered at the 2026 Global Technology, Media and Communications Conference on Tuesday, May 19, 2026 at 3:35 PM

Alexei Gogolev: OK, great. Hi, everyone, and welcome to J.P. Morgan Boston CMC Conference. My name is Alexei Gogolev, head of vertical SaaS team. Today, I'm delighted to be hosting one of the key founders and majority holder of Bentley Systems, company chairman, Greg Bentley. Greg, welcome.

First, Greg, I wanted to start by discussing the fact that the stock has traded as if the market is somewhat skeptical about Bentley's ability to fully capture upside from AI. In many of your conversations, you've noted that the market has discontinued and discounted your value in the near term.

What's your best case for why Bentley is positioned to be an AI leader and perhaps a long-term winner in this space?

Gregory Bentley: Alexei, the very long term is how we set our compass at Bentley Systems, and it perhaps is not surprising given our family majority ownership. That's the case with our generational succession in management. I was CEO for 30-plus years. Our new management team, new to their positions, but not new to the company necessarily, are in their 40s.

Their refrain is infrastructure AI as their opportunity. I had thought of that as rather audacious in the scheme of things, such a major opportunity as that. As things developed, I've come to be astonished at the pace at which AI is bringing the future forward for us, in what we have always wished to and aspired to bring about, better and more resilient infrastructure in the world.

That includes the operations and maintenance of infrastructure and opportunity, which is relatively newer for us. Of course, in bringing the future forward, it's the case that AI accelerates the pace of our own development on our roadmap to the future by multiplying the productivity of our software development.

I'm mainly talking about the inception of changes in the market structure for infrastructure that are truly auspicious for us. I've been talking this quarter about engineering firms, top design firms in

the world because they are, together, our largest constituency.

I was delighted this week to see the CEO of our largest account in his own comparable transcript say that, on the one hand, AI and technology's important to them because their backlogs are at record levels, as with everyone among our accounts.

They can't hire more people collectively in the industry because the number of infrastructure engineers only increases one percent per year. That can't change foreseeably, given the demographics of those retiring in comparison to the number in school, and that therefore, to do more work, they need to be more efficient. AI can help with that.

In response to an obvious question about whether AI will reduce their firm's work and revenues because their business model throughout infrastructure engineering is charging per hour for their work, with the efficiencies of AI, will they have lower revenue? He said, "I don't believe a bit of that. Our design work has not been optimized yet."

This is the opportunity we have been identifying with AI for engineering firms that in addition to the attended consumption where an engineer sits in front of our modeling and simulation applications, they will now, with agentic AI, be able to spin up agents that run hundreds and thousands of iterations of their potential designs to optimize in various dimensions and generate better and more valuable products for owner-operators.

He was talking about that value, the risk mitigation and so forth. In our own portfolio at Bentley Systems, these levels of optimization include across multiple disciplines. It includes optimizing the constructability of a project with our synchro applications.

It includes mitigating the risk of subsurface conditions, which is the principal risk that occurs in an infrastructure project.

Optimizing to reuse prior design modules to derisk and improve the quality. Then ultimately optimize the performance of the design by being informed by the performance of other designs for the same owner and how well they have performed in the field using our Bentley Infrastructure Cloud data about operations and maintenance.

All of these opportunities are coming together thanks to AI. When the design portion of an infrastructure project has broadly averaged about five percent of the total installed cost, and the portion of design billings spent on software has been about one percent of the design billings,

that's five basis points of the infrastructure project cost.

The infrastructure projects, which collectively across the world consume trillions of new dollars every year, five basis points of that is currently spent on software to make the design more intelligent and optimized and so forth.

The CEO's point was, when we're being paid by the hour, we only have time to do one version. If agents can help us, we'll be able to do a much better design. We'll be able to be paid for that. The owner will not be interested in optimizing and reducing the five basis points spent on software in the project when it can be a derisked project and perform much better.

Now over on the owner operator side, the opportunity for optimization and operations and maintenance is even larger.

What we call the digital twin opportunity is that the engineering logic that the owner pays for during design is at present, never used thereafter during the remaining operating life cycle of the asset, where 80 percent of the cost occurs in that maintenance. To be able to optimize maintenance would be able to do only the needed maintenance.

Our asset analytics opportunity uses drone capture of existing conditions, typically among other capture modalities, to have an as operated, calibrated model to which you then apply the engineering logic created during project design to make decisions about whether maintenance is needed and to be sure that the infrastructure remains safe.

Those opportunities accelerated by AI asset analytics is only made possible by AI have come together to create an economic opportunity for better and more resilient infrastructure in the world and get past the inhibitions and inertia that the existing institutional business model of charging by the hour have caused in the past.

At the same time as we have the technical capabilities from AI, we are getting past the commercial model obstacles to smarter and more resilient infrastructure. Much more will be spent on software and computing and infrastructure engineering.

I can't make the claim that that will all be to the benefit of Bentley Systems, but our position is ideal for this. In our 42 years, we are the trusted digital quartermaster for the engineering organizations and the owner operators, each are half of our business.

Our modeling and simulation applications, the ones you need to invoke agentically in order to optimize are the established standards in the world and trusted by the engineers.

Last of all, the data needed to do this optimization is in our project wise environment. It is the intellectual property of the engineering firm or the owner operator or both, but it positions us very well to help them themselves take advantage of that proprietary data in their proprietary differentiation to improve and make infrastructure better and more resilient.

Alexei: Greg, you released the STAAD MCP server and said you're exploring token-based pricing for API consumption. Can you give us an early data, how many accounts are actively using this STAAD MCP server today? What does a typical session look like in terms of compute intensity? When should investors expect the first commercially priced API consumption contract outside of asset analytics?

Gregory: It's only a couple of weeks now since this first of our MCP servers was released. Over the coming year, we'll completely outfit and instrument our application portfolio, both modeling and simulation applications. There's not a record yet. There are more than a couple of hundred downloads, and we don't think we can count those that are done through Claude itself.

What we can answer is the monetization question, which is that we'll have no monetization this year from API consumption or MCP consumption because our determination is to benefit the long term by having our accounts do as much exploration and experimentation with the APIs and MCP services as they can do with no intention to monetize it during this year.

We reckon that this year, we'll learn, for instance, about the costs, as to which you asked. The reason there's not a big risk to us in this year is that initially, the APIs and MCP services execute on the engineer's own desktop. That's a machine they have available to them anyway. They find that they can do hundreds of iterations at machine speed.

At the same time, they could have done one at human speed. We'll hopefully get experience as well in what the benefits are of these optimization approaches as we enable them to become richer and richer with more and more APIs across our modeling and simulation applications during the remainder of this year, Alexei.

Alexei: Thank you, Greg. At the most recent earnings call, you've laid out really well the math on representative, let's say, one million design project, and how software today is about $10,000 in that contract. You suggested that hypothetically doubling it to 20,000 with API consumption could

really save almost 20 percent of engineering labor.

That looks like a very compelling ROI, but what's the realistic adoption curve? Are any of your hundreds of accounts spending over a million dollars already and already running some agentic workflows in their production, or is this still entirely in pilot?

Gregory: We have 220 accounts that spend over a million dollars with us, 820 accounts that spend $250,000 or more, and these are big enough enterprises to have been already investing in AI R&D.

What they don't need to do with our agentic API road map is start at the bottom, replicating functionality that they couldn't replicate and wouldn't be able to validate anyway for modeling and simulation, but they can put that together into levels of optimization all the way up to the project level with the context they know about. That's a wonderful thing for them to be doing.

Many of them are already doing that, exercising our APIs that exist in our products today without our help. We don't have a way of tracking that. We don't have a way of monetizing it, but we want to encourage it anyway for the reasons that I just mentioned.

Alexei: In terms of certain bear case scenarios, it sounds like it could be that horizontal AI platforms with agents having access to open source structural analysis tools could eventually replicate some of that stat and what it does without paying Bentley for API access.

Would validated simulation engine and the regulatory requirement for proven tools, or maybe something else, prevent such a displacement?

Gregory: You're asking about our first MCP service is for STAAD. STAAD is our structural application that is the leader in the world in structural analysis. There are open-source structural analysis programs now that our accounts could be using.

The STAAD users are not. In general, they prefer the commercial products, which cost them only about three percent of an engineer's cost each day on average under our E365 program across our application set. There's not a great incentive to not need the best product set. STAAD understands the design codes, which the open source ones don't in their jurisdictions.

It does some optimizing at the element level, so its functionality is better. Generally, the engineering firms have much more to gain by spending their AI discretionary R&D investment on

leveraging the existing tools that are proven, that are mandated by their owner-clients, and which the engineers already know how to understand and use, and have chosen in these levels of optimization.

What I mean by levels of optimization, you could optimize material or cost for a structural solution, but a structural solution is only one of the disciplines that need to be solved in a design.

If you could orchestrate across those to have an objective function across the design, that would be good, but for a project, you might want to also optimize the availability of the materials, the type of labor, and maybe conformity with designs and assets the owner already uses, for instance.

These are all aspects of adding value that we can't do in the granular applications, and is the opportunity, preferably for both us and the account, for them to add, in optimizing at a higher level of value.

Alexei: How about the fact that you've highlighted that you explicitly need consent from your customers to train your own AI and those models? Obviously, of your competitors may consider using their customers' data without consent to train their models. Could your principled stance become a competitive handicap rather than an advantage longer term?

Gregory: Our principled stance is that our users' intellectual property is their intellectual property, and we're not entitled to appropriate it, to commingle it, to share it, especially not to have it help their competitors. That may seem different than other software companies' viewpoints on the matter, but our position is a bit different.

For one thing, before we get to the principle of it, consider the sensitivity of this data. These are the designs of the world's critical infrastructure. Consider whether anyone would want it uploaded to Claude or downloaded from ChatGPT, whatever, the vulnerabilities potentially for terrorists or adversary, etc., of the airports, and the wastewater, and grid, and so forth.

It can't be anonymized, by the way, because it's intrinsically geolocated, each such design. We have a responsibility to secure it and steward it. That doesn't disadvantage our account, whose data this is, because it's sufficient data for a given owner-operator of infrastructure. They're all large. There's lots of data that they have.

In the case of engineering firms, the same is the case. Our opportunity is to help them make the

most use of that data. An example would be, for instance, in automating drawing production, which is where 20 percent gets saved of the engineers' time. That allows them to work on an incremental project at a fixed pricing and improve the margins very substantially.

Automating drawing production can use the standards of a particular engineering firm, or perhaps a Department of Transportation would want all of its drawings to a particular dialect that works for them, so that the contractors to whom they're sending those drawings for construction bidding can expect a certain consistency.

That training can be done sufficiently with data which is the intellectual property of our account or which has agreed to share it. We do not even solicit permission from our accounts to have standing permission to access and use their data for AI training.

When we have particular projects to train models, we can ask and be authorized to access certain data that is not differentiated or confidential in the accounts judgment. We have a data agreement registry to keep track of that.

It will serve us very well because of the scale of these accounts and the amount of data they have to help them in a proprietary way, improve their differentiated offerings by using their own data, where we help them do that. That's the right compass setting for us, we think.

Alexei: That makes a lot of sense, Greg. Talking about one of the attractive businesses within your portfolio, Asset Analytics, it recently crossed $50 million revenue run rate. You've been clear that not all of it qualifies as ARR, because inspections aren't always annual.

Can you maybe help us size this more precisely longer term? What needs to happen to make most of this revenue annual? Is it some regulatory change that you're expecting, or maybe some pricing, packaging, or some technology change?

Gregory: To back up a bit, we have forever been promoting the potential of digital twins, which reuse the engineering intelligence developed for the project delivery in the first place, after delivery during operations and maintenance. In parts of the world like Asia, that's taken up right from the outset.

A new metro system will have a digital twin approach right from the start and use, and if it's done in Singapore, that will be the case, for instance. For owner operators, for instance, in the US, that's too complicated. They don't have many new projects starting. What asset analytics enables

is using new continuous surveying data capture technologies, drones in the main. I always do this wagging my finger to show a drone circling.

A drone produces video. Video is overlapping images. Our software resolves overlapping images into a 3D reality mesh that captures the digital twin of the operating asset today. Then you make it intelligent with AI recognizing everything and relating it to the engineering models.

That's the AI enables asset analytics to create an instant on digital twin, which you can use in operations and maintenance to know what operations and maintenance is required and to ultimately optimize that. It's a very powerful opportunity.

We started with cell towers and then roadway miles from dashcam data and Google Street View and now distribution poles, for instance, that together add up to this proven $50 million. It is recurring revenue, it's an annual subscription in each case, but it may be a subscription for a year when construction is occurring for a particular pole or tower or asset or mile.

It may not necessarily be renewed for subsequent years. In the case of distribution poles, an inspection is required every five years in the United States. Our hesitation in classifying it as ARR, it is recurring revenue. It may not be annually recurring if you don't renew it every year. It may be somebody different that you need to sell to after the construction or after the inspection.

We're working on that. The best way to work on that to make sure that the surveying by drone or otherwise is in fact continuous and occurs regularly is to make it simpler, have more AI, extract more value from it, and make it more worth doing more often, which in turn by the way if you knew you were going to inspect something regularly, you wouldn't need to replace it quite so often because you would know you're going to be able to catch a problem before it becomes significant, which the engineering model would tell you.

At any rate, you can imagine also in terms of the ARR challenge for an account having floors and ceilings because they may be doing different assets each year and so forth. We're doing all of those things to make it an ARR business substantively and contractually. The most important thing is to capture as many assets, get out there in front of this big opportunity that we can do.

Gregory: Greg, since acquiring a business called Seequent, you've grown subsurface ARR quite materially in civil infrastructure, I think you were saying by a factor of four. I don't know what dollar amount that equates to, but what's the addressable opportunity for that division? How many of your top 470 design firms accounts are already using Seequent for ground investigation?

Gregory: Sequent was what we call a platform acquisition almost five years ago now. It is a global company that started with 3D modeling below the ground for mining in particular. We saw the opportunity to apply that technology, which is different technology below the ground. You can't see and measure and take pictures.

You need to do boreholes and samples and drills and then infer mathematically what the strata are below the ground and where the water is seeping and environmental modeling of that. We saw the applicability of that to civil infrastructure projects where ground conditions are the biggest source of risk that cause project overruns.

The mining and resources business has also turned out to be good. The Sequent business has multiplied while we've owned it, but the civil portion of it has grown to be a very significant double-digit proportion. There's lots of headroom there as well.

It happens that not quite half of our top design firm accounts that you're referencing use Sequent at this point for subsurface modeling for ground-informed design, for derisking projects, and they all should for all their projects. We'll work continue to work on that opportunity.

Alexei: Perfect. Switching to E365, I understand it's now 46 percent of ARR and growing pretty rapidly. Can you walk us through the unit economics of E365 renewals? What's the typical floor uplift you're achieving at those renewals? How does that compare to maybe a few years back?

Gregory: E365 is a pure consumption program. We're paid for, and our revenue and our ARR are strictly based on, for our open applications, the number of days used for our project-wise environment, which is 20 percent or 30 percent of it. It's per quarter in that case.

The great aspect of this is there's no multi-year contracting and commitment or accounting. It's strictly ratable revenue recognition back to the good old days before 06/06. What occurs at renewal is an annual price escalation becomes effective.

There tends to be some increase in ARR, but we were asked by accounts, although this is not an aspect of the program contractually, if we would grant them a ceiling on what their consumption could become in the ensuing year. We said, "Yes, if you'll grant us a floor." A collar became conventional for the contract to be negotiated a contract renewal.

Then, a few years ago, accounts who saw this sort of thing elsewhere said, "Can we negotiate

that for multiple years in the future so that we'll have a cap on what we could potentially spend based on how our consumption goes over time?" We said, "Yes, if we could symmetrically negotiate a floor as well."

That has become the norm for each of these renewal negotiations. What happens at renewal mechanically is our ARR would be increased by the price escalation, which tends to be in mid-single digits each year.

The remainder of the consumption increase is in the main there upsell to use more specialized applications of ours, a better, for us, mix of product usage. Not much of it anymore comes from increased volume because there is a limited number infrastructure engineers, as we said earlier. There's information in the result of this negotiation.

I always ask each quarter, "Where did we wind up with the average of the negotiated collar increase escalation out into the future?" The answer for as long as we've been doing it is 10 percent or so on average across the world and across all account sizes and so forth, which is the mentality.

We can gauge from that that infrastructure engineering organizations, the substantial ones. These are the ones the two-thirds of that pay us $250,000 a year and more, are confident in their need for and expectation for. What they'll actually pay us is what they actually consume.

It's the most fair contract format you can imagine, but they're confident enough about their next several years, notwithstanding AI concerns and threats and so forth, to be able to commit to floors, and we grant ceilings at that 10 percent level. That provides a lot of consistency in our ARR, it must be said.

I take the view as an owner of the business that it's so much consistency, it allows us to take on some risk on the margin, like the asset analytics business, which is less predictable, going after big contracts and so forth. We're to have a happy medium while we work on building that one up, ultimately to be comparable.

In the meantime, AI is on everyone's mind, and it's not a deterrent to continuing to commit to the level of attended consumption because that's what we're charging for now for our existing applications.

Alexei: Perfect, Greg. Final question. You've talked a lot about having capacity for up to $400

million worth of annual programmatic acquisitions without increasing leverage. Last year was a little light, so you spent less than $100 million on acquisitions.

What do you feel is holding you back from deploying more capital aggressively? Is it lack of targets, maybe some valuation discipline, or integration bandwidth? Now that you've expanded beyond asset analytics, what other domains are you looking at?

Gregory: For a period of time after our platform acquisitions, that's the billion-dollar-scale acquisitions of Sequent, which was at the time about 10 percent of our size, and then PowerLine Systems the next year, we issued convertible bonds that, if you consider them debt, brought our leverage ratio up to 5X or so.

Over several years, we've worked that down now to be under 2X, which we regard to be tolerably optimal. We have no wish to delever further, and that gives us the number, given our cash flow is above $500 million last year and this coming year to be able to do $400 million or so of acquisitions. That is our priority.

We would like to especially acquire AI for additional asset types and asset analytics beyond the ones I mentioned just to get the time advantage of not needing to take the time to do the training ourselves. They're scarce to find those good assets, and it took us until the end of last year to do the couple of acquisitions we did do while that was our priority.

Now, while we've reached an optimal leverage, we'll continue also to repurchase our shares, but the greater priority will be such acquisitions, including now, since we can expand beyond the asset analytics opportunity, where we are disciplined.

It's not the case that we're allergic to dilution in margins compared to our very good margins, but we are disciplined in valuation compared to private equity firms, perhaps, still.

We are now returning to prioritize as well acquisitions of established simulation engines, especially that have the value in this agentic immediate future that I'm describing, of allowing us to optimize on even more dimensions of engineering performance. You'll see, hopefully, both of those. We really do hope so in this year.

Alexei: Perfect. Greg, thank you very much for joining us today. Appreciate your time.

Gregory: Cheers. Thank you.

[applause]