


Second Life creator Linden Lab starts testing its virtual-reality world: Project Sansar

Dean Takahashi

@deantak

August 18, 2015 8:00 AM



Linden Lab is prepping Project Sansar, a VR world.

Image Credit: Linden Lab

Connect with top gaming leaders in Los Angeles at GamesBeat Summit 2023 this May 22-23. [Register here.](#)

Second Life creator [Linden Lab](#) is announcing today that it has begun inviting creative people to test its new virtual reality world, code-named Project Sansar.

Project Sansar should debut to the general public in 2016, and its aim is to democratize virtual reality as a creative medium. The platform aims to empower people to easily create, share, and

<https://venturebeat.com/games/linden-lab-starts-testing-its-virtual-reality-world-project-sansar/>

1/15

monetize their own multi-user, interactive virtual experiences, without requiring engineering resources.

Linden Lab said the platform will enable professional-level quality and performance with “exceptional visual fidelity, 3D audio, and physics simulation.”

Project Sansar will be optimized for VR headsets like the Oculus Rift, but also accessible via PCs and (at consumer launch) mobile devices.



EVENT

GamesBeat Summit 2023

Join the GamesBeat community in Los Angeles this May 22-23. You'll hear from the brightest minds within the gaming industry to share their updates on the latest developments.

[Register Here](#)

Linden Lab was founded in 1999 and it triggered huge interest in virtual worlds. After a big shakeout, Linden Lab's Second Life was one of the survivors. But its technology looks a bit backward today, and Project Sansar is an attempt to return to the high ground of entertaining virtual experiences.

Project Sansar will reportedly have exploration, socializing, and other experiences made possible through expressive avatars, using text, and voice chat. And it will make it easy and fun for users to create their own social VR experiences, eliminating the need to find a professional developer to build such an experience today.

Linden Lab said users won't have to worry about issues such as hosting and distribution, multiuser access and communication systems, virtual currency, and regulatory compliance.

“We want to lower the barrier of entry for VR experience creation,” said Linden Lab CEO Ebbe Altberg in a statement. “Project Sansar will do for virtual experiences what WordPress has done for the web: empower a broad range of people to create with professional quality and reach global audiences.

“By greatly expanding who can create virtual experiences, Project Sansar will also extend the value of VR to a wide variety of use-cases – from gaming and entertainment to education, architecture, art, community-building, business meetings, healthcare, conferences, training, and more.”

At the outset, the creators invited to help test Project Sansar will create 3D content using Autodesk's Maya software. At the consumer launch next year, Project Sansar will support a variety of third-party creation tools as well as native building options, so that creators can work with their preferred software.

GamesBeat's creed when covering the game industry is "where passion meets business." What does this mean? We want to tell you how the news matters to you -- not just as a decision-maker at a game studio, but also as a fan of games. Whether you read our articles, listen to our podcasts, or watch our videos, GamesBeat will help you learn about the industry and enjoy engaging with it. [Discover our Briefings.](#)

Find your dream job on VentureBeat Careers

**IBM**

Armonk

View 321 Jobs

<https://jobs.venturebeat.com/company/ibm-2>**Apple**

Cupertino

View 1221 Jobs

<https://jobs.venturebeat.com/company/apple-2>**GE Healthcare**

Chicago

View 8 Jobs

<https://jobs.venturebeat.com/company/ge-healthcare-2>

**Samsung Electronics**

Suwon-si

View 69 Jobs


<https://jobs.venturebeat.com/company/samsung-electronics-2>

VB Insight

Can healthcare show the way forward for scaling AI?

VB Staff

March 15, 2023 09:00 AM



This article is part of a VB Lab Insights series on AI sponsored by Microsoft and Nvidia.

Don't miss additional articles in this series providing new industry insights, trends and analysis on how AI is transforming organizations. [Find them all here](#)

Scaling [artificial intelligence \(AI\)](#) is tough in any industry. And healthcare ranks among the toughest, thanks to highly complex applications, scattered stakeholder networks, stringent licensing and regulations, data privacy and security — and the life-and-death nature of the industry.

“If you mis-forecast an inventory level because your AI doesn’t work, that’s not great, but you’ll recover,” says Peter Durlach, Executive Vice President and Chief Strategy Officer of [Nuance Communications](#), a conversational AI company specializing in healthcare. “If your clinical AI makes a mistake, like missing a cancerous nodule on an X-ray, that can have more serious consequences.”

Even with the current willingness of many organizations to fund AI initiatives, many healthcare organizations lack the skilled staff, technical know-how and bandwidth to deploy and scale AI into clinical workflows. In fact, it’s far lower than the [average](#) of around 54% for all industries combined.

Despite the difficulties, machine learning (ML) and other forms of AI have impacted a wide range of clinical domains and use cases in hospitals, R&D centers, laboratories and diagnostic centers. In particular, deep learning and computer vision have helped improve accuracy, accelerate interpretation and reduce repetition for radiologists for x-ray, CT, MR, 3D ultrasound and other imaging. With global shortages of [radiologists](#) and [physicians](#) looming, AI assistance could be a “game-changer.”

After slow growth that has [trailed nearly every industry](#), many analysts forecast that healthcare AI will boom in 2023 and beyond. The global market is expected to [exceed \\$187 billion](#) by 2030, reflecting fast-growing demand.

To take advantage of investments, enterprises and industry vendors must overcome several technical obstacles to adoption of clinical AI. Chief among them: Lack of standardized, healthcare-specific platforms and integrated development and run-time environments (IDEs and RTEs).

Moreover, current infrastructure often lacks the functionality, workflows and governance to easily create, validate, deploy, monitor and scale — up, down and out. That makes it difficult to scale up during a morning clinic, then scale down during the evening when demand is lower, for example. Or to easily expand deployment of AI systems and models across organizations.

Yet despite (and perhaps because of) these challenges, some of today’s most innovative and effective approaches for moving AI into production come from healthcare.

What follows are conversations VB had separately with two global leaders about leading-edge, cloud-based approaches that might offer blueprints for other industries struggling with scaling automation.

1. Nuance: 'From Bench to Bedside,' deploying for impact

Accelerating creation and deployment of trained models at scale with a secure cloud network service — a conversation with Peter Durlach, Executive Vice President and Chief Strategy Officer at Nuance.

Good news: The growing popularity of foundation and large language approaches is making it easier to create AI models, says Durlach. But the difficulty of deploying and scaling AI models and applications into healthcare workflows models continues to present a formidable challenge.



Credit: Nuance

“About 95% of all models built in-house or by commercial vendors never get deployed, because getting them into clinical workflow is impossible,” Durlach said. “If I’m a client building a model just for myself, it’s one set of challenges to get that deployed in my own company. But if I’m a commercial vendor trying to deploy across multiple settings, it’s a nightmare to integrate from the outside.”

Making it easier for hospitals, AI developers and others to overcome these obstacles is the goal of a new partnership between Nuance, Nvidia and Microsoft. The aim is to [simplify and speed](#) the translation of trained AI imaging models into deployable clinical applications at scale by combining the nationwide [Nuance Precision Imaging Network](#), an AI-powered Azure cloud platform, and [MONAI](#), an open-source and domain-specialized medical-imaging AI framework cofounded and accelerated by Nvidia.

The latest solution builds on two decades of work by Burlington, Mass.-based Nuance to deploy AI applications at scale. “We are a commercial AI company,” Durlach explains. “If it doesn’t scale, it has no value.” In these interview highlights, he explains the value of an AI development and deployment service and suggests what to look for in a provider of AI delivery networks and cloud infrastructure.

Underestimating complexity

“People underestimate the complexity of closing the gap from development to deployment to where people actually use the AI application. They think, *I put a website up, I have my model, I have a mobile app*. Not so much. The activities involved in implementing an AI stretch from R&D through deployment to after-market monitoring and maintenance. In life science, they talk about getting a clinical invention from the bench to the bedside. This is a similar problem.”

Key steps in developing and using AI for medical imaging



Credit: Nuance


The value of specialized cloud-based delivery and development

“If I’m a healthcare organization, I want to use AI to drive very specific outcomes. I don’t want to have to build anything. I just want to deploy an application that solves a specific business problem. Nuance is bringing end-to-end development, from low-level infrastructure and AI tools all the way up to specific deployable applications, so you don’t have to stitch components together or build anything on top.

“The Nuance Precision Imaging Network runs on Azure and is accessible across more than 12,000 connected facilities across the country. A health system or a commercial vendor can deploy from development to runtime with a single click and be already integrated with 80% percent of the infrastructure in U.S. hospital systems today.

The new partnership with Nvidia brings specialized ML development frameworks for medical imaging into clinical translation workflows for the first time, which really accelerates innovation and clinical impact. [Mass General Brigham](#) is one of the first major medical centers to use the new

offering. They're defining a unique workflow that links medical-imaging model development, application packaging, deployment and clinical feedback for model refinement."



Choosing a cloud infrastructure vendor

"When Nuance was looking for cloud and AI in healthcare, one of the first things we asked was *What's the company's stance on data security and privacy?* What are they going to do with the data? The large cloud companies are all great. But if you look closely, there are many questions about what's going to happen to the data. One's core business is monetizing data in various ways. Another one often uses data to go up the stack and compete with their partners and clients.

"On the technical side, each cloud company has their strengths and weaknesses. If you look at the breadth of the infrastructure, Microsoft is basically a developer platform company that provides tools and resources to third parties to build solutions on top of. They're not a search company. They're not a pure infrastructure company or a retail company. For us, they have a whole set of tools — Azure, Azure ML, a bunch of governance models — and all the development environments around .NET, Visual Studio, and all these things that make it easier, not trivial, to build and deploy AI products. Once you're running, you need to look closely at scalability, reliability and global footprint.

"For data security, privacy and comfort with the business model, Microsoft stood out for us. Those were major differentiators.

“Nuance was acquired by Microsoft about 10 months ago. But we were a customer long before that for all these reasons. We continue running and building atop Microsoft, both on-premises and in Azure, with a wide array of Nvidia GPU infrastructure for optimized training and model building.”

Focus on value, not technology

“AI technology is only as good as the value it creates. The value it creates is only tied to the impact it drives. The impact only happens if it gets deployed and adopted by the users. Great technical people look at the end-to-end workflow and the metrics.

“Don’t get lost in the technology weeds. Don’t just get caught up in looking at one tool set or one annotation tool or one inferencing thing. Instead, ask *What is the use case? What are the metrics that the use case is trying to move around cost, revenue? What is required to actually get the model deployed?* Getting super rigorous around that and not underestimating and falling in love with building the model. It has almost no value if it doesn’t end up in the workflow and drive impact.”

Bottom line: Taking advantage of an established commercial delivery network and cloud ecosystem lets you focus on developing and refining AI models and applications that deliver clear value and help drive key organizational goals. When choosing a network and cloud provider, look closely at three key areas: how their business models impact data privacy, the completeness of their AI development and delivery environment, and their ability to easily scale as widely as you require.

2. Elekta: Collaborate to ‘dream bigger’ and speed innovation of products and AI

Scaling global R&D infrastructure in the cloud helps make next-gen, AI-powered radiation therapy more accessible and personalized — a conversation with Rui Lopes, Director of New Technology Assessment at [Elekta](#).

In 2017, Rui Lopes visited a major radiology conference and noticed a big change. Instead of “big iron and big software,” which usually took up most of the floor space, almost half of the trade show was now dedicated to AI. To Lopes, the potential value of AI for cancer diagnosis and for cancer treatment was undeniably clear.

“For clinicians, AI offers an opportunity to spend more time with a patient, to be more care-centric rather than just being the person in the darkroom who looks at a radiograph and tries to figure out if there’s a disease or not,” says Lopes, Director of New Technology Assessment for Elekta, a global

innovator of precision radiation therapy devices. “But when you recognize that a computer can eventually do that better at a pixel scale, the physician starts to question, what is my real value in this operation?”

Today, the growing openness of healthcare professionals worldwide to asking that question and to embrace the opportunity of cancer care driven by AI is due in no small part to Elekta. Founded in 1972 by a Swedish neurosurgeon, the company gained international renown for its revolutionary [Gamma Knife](#) used in non-invasive radiosurgery for brain disorders and cancer, and most recently its groundbreaking Unity integrated MR and linac (linear accelerator) device.

For much of the last decade, Elekta has been developing and commercializing ML-powered systems for radiology and radiation therapy. Recently, the Stockholm-based company even created a dedicated radiotherapy AI center in Amsterdam called the POP-AART lab. The company is focusing on harnessing the power of AI to provide more advanced and personalized radiation treatments that can be quickly adapted to accommodate any change in the patient during cancer treatments.



[Credit: Elekta](#)

At the same time, Elekta recently launched its “Access 2025” initiative that aims to increase radiotherapy access by 20% worldwide, including in underserved regions. Elekta hopes that by integrating more intelligence into their systems they can help overcome common treatment

bottlenecks such as shortages of clinician time, equipment and trained operators, and as a result, ease the strain on patients and healthcare providers.

Along the way, Elekta has learned valuable lessons about AI and scaling, Lopes says, even as company expertise and practices continue to evolve. In these interview highlights, Lopes shares his experience and key learnings about [moving to on-demand cloud infrastructure and services](#).

Wanted: Smarter collaboration and data sharing

“We’re a global organization, 4,700 employees in over 120 countries, with R&D centers spread across more than a dozen regional hubs. Each center might have a different priority for improving a particular product or business line. These disparate groups all do great work, but traditionally they each did it in a bit of isolation.

“As we considered how to ramp up the speed of our AI innovations, we recognized that a common scalable data infrastructure was key to increasing collaboration across teams. That meant understanding data pipelines and how to manage data in a secure and distributed fashion. We also had to understand the development and operational environment for machine learning and AI activities, and how to scale that.”

Costly on-premises servers, ‘small puddles of data’

“As a company, we have traditionally been very physics-based in our research in radiotherapy. Our data and research scientists were all very on-prem-centric for data management and compute. We invested in large servers through large capital purchases and did data preparation and massaging and other work on these local machines.

“AI has a voracious appetite for data, but because of privacy concerns, it’s a challenge to get access to large volumes of medical data and medical equipment data required to drive AI development. Luckily, we have very good, very precious partner research relationships around the world, and we employ different techniques to respect and maintain strict privacy requirements. But typically, these were small puddles of data being used to try to drive AI initiatives, which is not the ideal formula.

“One thing we did early was establish a larger-scale pipeline of anonymized medical data that we could use to drive some of these activities. We didn’t want replication of this data lake across all our distributed global research centers. That would mean people would have different copies and different ways of managing, accessing and potentially even securing this data, which we wanted to keep consistent across the organization. Not to mention that we’d be paying for duplicate

infrastructure for no reason. So, a very big part of the AI infrastructure puzzle for us was the warehousing and the management of data.”

Budgeting mind shift: Focus on cloud’s new capabilities

“As we delved more and more into ML and AI, we evaluated the shift from on-prem compute to cloud compute. You do a couple of back-of-the-envelope calculations first: Where are you regionally? What are you paying now? What type of GPUs are you using? As you’re starting this journey, you’re not quite sure what you’re going to do. You’re basing the decision on your current internal capacity, and what it would cost to replicate that in the cloud. Almost invariably, you end up thinking the cloud is more expensive.

“You need to take a step back and shift your perspective on the problem to realize that it’s only more expensive if I use [cloud] the way I use my on-prem capacity today. If instead you consider the things you can do in cloud that you can’t do onsite – like run parallel experiments and multiple scenarios at the same time or scale GPU capacity – the calculus is different. It really is a mind shift you have to make.

“As you think of growth, it becomes obvious that migrating to cloud infrastructure can be extremely advantageous. Like with any migration, you have a learning curve to becoming efficient and managing that infrastructure properly. We may have forgotten to ‘turn off the lights’ on capacity a couple of times, but you learn to automate much of the management as well.”

Aha moment: Leverage smart partners

“I mentioned the challenges of accessing medical data. But another part of the challenge is that often the data you need to access is a mix of types and standards or consists of proprietary formats that can change over time. You want any infrastructure you build to have flexibility and growth capabilities to accommodate this.

“When we looked around, there was no off-the-shelf product for this, which was surprising and a big ‘aha moment’ for us. We quickly recognized this was not a core competence for us – you really need to work with trusted partners to build, design and scale out to the right level.

“We were fortunate to have a global partnership with Microsoft, who really helped us understand how best to create an infrastructure and design it for future scaling. One that would let us internally catalog data the right way, allow our researchers to peruse and select data they needed for developing AI-based solutions – all in a way that is consistent with the access speed and

latency we were expecting, and the distributed nature of our worldwide research teams and our security policies.”

Starting smart and small

“We started limited pilots around 2018 and 2019. Rather than betting the bank on a massive and ambitious project, we started small. We continued our current activities and way of working with the on-premises and non-scalable systems, setting aside a little bit of capacity to do limited experiments and pilots.

“Setting up a small Azure environment allowed us to create virtual compute as well and doing a redundant run of a smaller experiment and asking ‘What was that experience?’ This meant getting faster, more frequent small wins instead of risking large-project-fatigue with no short-term tangible benefits. These, in turn, provided the confidence to migrate more and more of our AI activities to the cloud.

“With COVID and everybody holed up at home, the distributed virtual Azure environment was very practical with a level of facility and convenience we didn’t have before.”

Learning new ways and discipline


“We recognized that we needed to learn as an organization before really jumping into [cloud-based AI]. Learning from it, too, so that parts of the team were getting exposed, understanding how to operate in the environment, how to use and properly leverage the virtual compute capacity. There’s operational and knowledge inertia to overcome. People say: ‘There’s my server. That’s my data.’ You have to bring them over to a new way of doing things.

“Now, we’re in a different space, where the opportunity is much bigger. You can dream bigger in terms of the scale of the experiment that you might want to do. You might be tempted to try to run a really large learning on a massive dataset or a more complex model. But you must have a bit of discipline, walking before you run.”

Help wanted: Developing new products, not models

“Rather than going out and recruiting boatloads of AI experts and throwing them in there and hoping for the best, we recognized we needed a mix of people with domain knowledge of the physics and radiotherapy.

“We did a few experiments where we brought in some real hardcore AI people. Great people, but they’re interested in developing the next great model architecture, while we’re more interested in applying solid architectures to create products to treat patients. For us, the application is more important than the novelty of the tech. At least for now, we feel there has to be organic growth, rather than trying to throw an entire new organization or a new research group at the problem. But it’s a challenge; we’re still in the process.”



IT as trusted partner and guide

“I’m in the R&D department, but we interact with the IT department very closely. We interact with the sales and commercial side very closely, too. Our Head of Cloud, Adam Moore, and I have more and more discussions about sharing learnings across corporate initiatives, including data management and strategy and cloud. Those are strands of the DNA of the company that are going to be intertwined as we move forward, that will keep in lockstep.

“If you’re lucky, IT is a red thread that can help through all of that. But that’s not always the case for many companies or entire IT departments. There’s a competence buildup that needs to happen within an organization, and a maturity level within IT. They’re the sherpa on this journey that hopefully helps you get to the summit. The better the partner, the better the experience.”

Toward more universal treatment and ‘adaptation’

“More centers and physicians are embracing the belief that (AI-assisted radiology) can have a positive impact and allow them to get closer to what’s most important — providing the best, personalized care to patients that’s more than just cookie-cutter care because there’s no time to do anything but that.

“AI is not only helping with the productivity bottleneck, but with what we call [adaptation](#). Even while the patient is on the table, about to be treated, we can take clinical decisions and [dynamically modify things on the fly](#) with really fast algorithms. It can make these hour- or day-long processes happen in minutes. It’s beyond personalization, and it’s really exciting.”

Bottom line: Focus early on data pipelines and infrastructure. Start small, with smart partners and close partnership between IT and the groups developing AI. Don’t get sidetracked by “apples- to-oranges” cost comparisons between cloud and on-premises environments. Instead, expand your vision to include new capabilities like on-demand parallel processing and HPC. And be prepared to patiently overcome organizational inertia and build up new competencies and attitudes toward data sharing.

VB Lab Insights content is created in collaboration with a company that is either paying for the post or has a business relationship with VentureBeat, and they’re always clearly marked. For more information, contact sales@venturebeat.com.

GamesBeat

[Press Releases](#) [Contact Us](#) [Advertise](#) [Share a News Tip](#)

[Contribute to DataDecisionMakers](#)

[Careers](#) [Privacy Policy](#) [Terms of Service](#)
[Do Not Sell My Personal Information](#)

© 2023 VentureBeat. All rights reserved.