The NΞXXT Perspective

Digital signage has long lived at the intersection of content, engagement, and data. It's one of the few domains in ProAV with a legacy of capturing real-time audience signals—long before AI became the industry's headline. As such, it has become a proving ground for applied intelligence: measuring gaze, reacting to presence, tailoring media in real time. These aren't theoretical futures—they're deployed realities.

In this personal and technically rich reflection, NΞXXT contributor Greg Herlein gives us a front-row seat into what happens when those capabilities evolve. When embedded NPUs and open-source models make edge AI not only possible, but practical. His story is a microcosm of the broader shift: from centralized to distributed, from locked systems to open platforms, from specification to potential, and from guesswork to insight.

But I'd argue that this isn't just a Digital Signage story. The lessons Greg draws - from privacy-first design to community-driven integration have the capacity to be inspirations, even blueprints, for the rest of the AV and collaboration landscape. What if the same tools used to measure audience engagement at a retail screen could help us understand real impact in a boardroom or classroom? What if edge agents could not only play content but track performance, detect anomalies, or even suggest better outcomes?

As we move from experimentation to enablement, the innovation at the edge isn't peripheral—it's foundational. And thanks to leaders like Greg, we're beginning to understand how building for the edge means building for what's NΞXXT.

NEXXT
TRANSFORMATION STARTS NOW

A Personal Journey into AI at the Edge

By Greg Herlein - NEXXT Contributor & VP Software at BrightSign

I've spent over 15 years in the digital signage industry, both as an early signage technology developer/operator and now as a technology leader. I've watched many transformative technologies arrive, like 4K and HTML5/webapps, then 8K. None of those are as exciting as the integration of AI capabilities directly into signage players.

This isn't just about jumping on the AI bandwagon - it's about solving real problems that I've personally encountered throughout my career in ways that weren't previously possible or cost-effective.

One early solution that is an example of "AI at the Edge" is a camera-based "machine vision" application to analyze and react to the audience. Our industry often wants to "measure how many people saw the media" or "we want to know when to play a specific ad based on who is present at the screen." Solutions for this have been available for some time, but often not cost-effective or practical from a deployment perspective.

Many existing solutions require a connection to the cloud to achieve the needed performance. Those solutions often create privacy and cost problems. Doing the machine vision directly on the player saves costs and has the benefit that the images never leave the player, alleviating many privacy concerns. The breakthrough that makes this kind of solution possible is embedded in new media players: the Neural Processing Unit (NPU). These are like GPUs, which are for graphics, but instead are optimized to run AI models. It's like a math-coprocessor for AI.

So we set out to prepare an example that uses an inexpensive USB camera and the onboard NPU to run an AI model to analyze what the camera "sees" and to measure the audience and then react. This article is about that journey.

The Power of Open Integration

Over my years in this industry, I've learned that the most successful solutions are built on platforms that embrace integration at every level of the stack. When I was deploying large retail networks a decade ago, I specifically chose players that offered GPIO and serial control on the hardware side, and robust APIs and scripting interfaces on the software side. This openness allowed us to easily integrate into our content management software and to enable features far beyond basic playback. Today many customers demand such an open approach.

One exciting aspect of today's AI revolution is the same spirit of openness and enablement. Several of the better Large Language Models (LLMs) are completely open source, and all of them have open APIs. These enable AI use in two ways: sophisticated developers can integrate with code at the API level, but end-users can leverage the models immediately for practical use.

I believe this dual approach is powerful. The rapid advancement of open-source AI models commoditizes many applications. At the same time, companies with specialized expertise can bring their refined solutions to the edge. For example, Quividi has spent many years fine-tuning their real-time audience analysis solution in retail.

The Power of Community and Open Source

Throughout my career, I've seen how open source and community-driven development can accelerate innovation. What's particularly exciting in this new wave of AI integration is seeing companies embrace this philosophy. For example, at BrightSign we recently created a "gaze detection" application using openly available AI models. The application sends simple network messages about audience engagement that any content management system can understand. This kind of open approach means that anyone can build on these foundations to create their own innovative solutions. I'm especially encouraged by companies releasing their AI integration examples under open-source licenses. The audience measurement app that we built at BrightSign is fully open sourced. This kind of sharing not only helps individual developers but strengthens the entire digital signage ecosystem. When companies provide documented, working examples of AI integration, it gives everyone a foundation to build upon, dramatically reducing the time and effort needed to bring new solutions to market.

Why Edge AI Matters

We are experiencing a step-change moment in how humans interact with computers. AI is enabling tremendous changes in productivity and capability. But not every application can connect to the cloud. Not only is that simply not cost-effective, it's likely to be too slow for many applications and raise substantial privacy worries. In digital signage, doing this work at the edge, directly on the player, creates many possible use cases.

Imagine a voice interactive kiosk, for example, where the retail customer can ask questions about the product on display. Doing that work completely on the player opens vast new possibilities for one-to-one marketing and shopper education.

Imagine having a smart agent on every player that is constantly evaluating media playback, metrics and logs, watching for anomalies and making intelligent status reports. If you are operating tens of thousands of players that's a huge gain for your operational costs and overall system uptime.

These are only two possible use cases for Edge AI in our industry. The ability to run AI models at the edge, without any extra hardware cost beyond what you already pay for the player, changes everything. We've only just begun to explore what we can do with this technology.

Looking Forward

I'm incredibly optimistic about the future of AI in digital signage. The combination of edge AI capabilities, open-source collaboration, and industry expertise is creating unprecedented opportunities for innovation. The key will be maintaining our focus on solving real problems while fostering an open, collaborative community that can turn these technological possibilities into practical solutions.

What excites me most is seeing how different players in our ecosystem - from individual developers to large companies - are coming together to push the boundaries of what's possible. By sharing knowledge, code, and experiences, we're not just advancing individual projects; we're elevating the entire industry. I look forward to being part of this journey and seeing the innovative solutions our community will create.