THE WORKSTATION AS THE NEW AI FRONTIER

Sponsored by

DELL Technologies

NVIDIA

Produced by

GEO WEEK NEWS

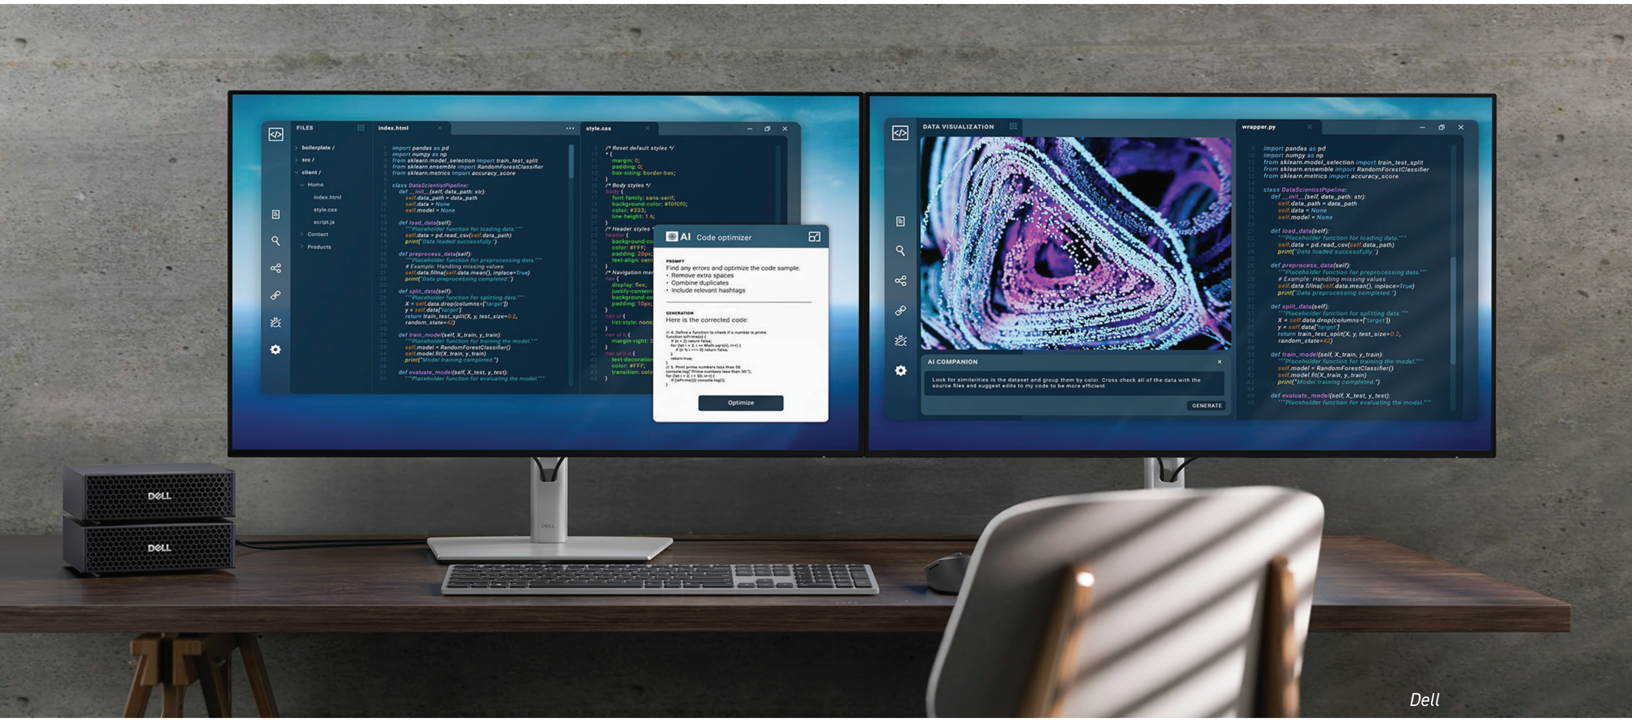

Dell

Introduction

The geospatial industry is moving at unprecedented speed. AI-powered software features are arriving at an accelerating pace, and the workflows that they enable are transforming what is possible in the field and on the job site. Tools for automated point cloud segmentation, real-time object classification, and other AI-driven workflows are already in use, the datasets are getting bigger, the tools are getting smarter, and the expectation for speed and accuracy has never been higher.

For many teams, the hardware is lagging behind the software that depends on it. Most workstations in use today were built for a different era, one where processor speeds and memory were the metrics that mattered. Those specs still matter, but they no longer tell the whole story. The workflows driving the industry forward now run on accelerated and parallel compute, and that requires a fundamentally different type of machine.

There is also a broader assumption worth questioning. Over the last decade, as cloud computing became more accessible, the idea took hold that heavy processing belonged somewhere else, offloaded, remote, and on-demand. For some use cases that model works, but for geospatial professionals working in the field with massive datasets, limited bandwidth, and connectivity that cannot be guaranteed, it introduces risks that are difficult to ignore. Moving that volume of data to and from the cloud also carries real egress costs that accumulate quickly and are easy to underestimate.

The answer is not pushing data further out into the cloud. It is bringing capable compute closer to where the data lives. A new generation of workstations built around accelerated and parallel compute is putting serious AI capability directly at the point of work, and the industry is taking notice.

The workstation is back, and needs to be at the center of the AI conversation.

The Evolution of the Workstation

For a long time, workstation performance was a linear conversation. You looked at clock speed, RAM, storage throughput. Faster processor, better machine. That framework made sense when the work was primarily sequential, one task at a time, processed in order. However, AI does not work this way.

The workflows reshaping geospatial today, AI-assisted point cloud segmentation, automated feature detection, real-time lidar processing, and radiance field reconstruction, are massively parallel workloads that require thousands of operations running simultaneously. That is a fundamentally different kind of computing, and it calls for a fundamentally different kind of hardware.

As agentic AI tools become part of professional workflows, the demands go further still. Local accelerated compute can act as an AI offload engine for task-specific agents, classifying site data, preparing project summaries, flagging anomalies, generating design alternatives, and coordinating multi-step analysis across local applications and datasets. The workstation is no longer just running the work. In an agentic AI model, it is increasingly directing it.

The shift is from processing power to accelerated compute. And the engine driving that shift is the GPU.

For most of its history, the GPU lived in the workstation as a specialized tool, handling graphics, rendering, and visualization. What has changed is the recognition that the same architecture making GPUs exceptional at rendering, the ability to execute enormous numbers of smaller calculations in parallel, makes them equally powerful for AI inference, deep learning, and the kind of intensive data processing that modern geospatial workflows demand.

As agentic AI becomes part of everyday professional workflows, that GPU headroom becomes even more critical. Running multiple task-specific agents concurrently alongside the industry software applications they are coordinating with places significant parallel processing demands on the system. The GPU is no longer a supporting player. For AI-driven work, and increasingly for the agents orchestrating that work across multiple applications and datasets simultaneously, it is the primary engine.

Dell Technologies and NVIDIA have built their partnership around exactly this transition. Rather than leaving configuration and compatibility to chance, Dell Pro Max and Precision workstations with the latest NVIDIA RTX PRO GPUs are engineered, tested, and certified together, purpose-built for the accelerated compute demands of modern AI workflows. For geospatial teams, that means less time worrying about whether the hardware will keep up and more time focused on what the work can now produce.

Two Ways of Computing

	Traditional Compute (CPU)	Accelerated Compute (GPU)
How it works	Tasks executed sequentially, one after another	Thousands of operations executed simultaneously
Optimized for	Complex instructions, general purpose tasks	Parallel workloads, AI inference, deep learning
Geospatial example	Running a CAD file, managing project data	Processing a point cloud running object detection, generating a 3D mesh
The limitation	One task at a time means AI workflows slow to a crawl	Requires modernGPU-equipped hardware to unlock
The opportunity	Most teams already have a GPU in their workstation	That machine may already be a parallel compute device ready for AI

Enabling Disruptive AI and Reality Capture Technologies

The pace of change in AI-powered geospatial software is not slowing down. It is accelerating. New techniques and tools are arriving faster than most teams can fully evaluate them, and the gap between what the software can do and what the hardware running it can handle is widening in real time.

Two areas in particular illustrate how significant that shift is becoming.

Neural Radiance Fields and the more recently developed Gaussian Splatting technique, have captured serious attention across the geospatial and AEC communities.

The ability to reconstruct photorealistic 3D scenes from fewer images than traditional photogrammetry requires, with results that can be geo-located, measured, merged with point clouds, and rendered in real time, represents a meaningful leap forward. Generating and rendering these scenes in real time is computationally intensive work, and it is exactly the kind of massively parallel workload that modern GPU architecture was built to handle.

Point cloud segmentation tells a similar story.

Tools like Leica Cyclone HDR are leveraging CUDA cores on NVIDIA GPUs to automate object detection and classification tasks that once required painstaking manual effort. The parallel processing architecture that makes this possible is the same architecture at the heart of modern Dell Pro Max and Precision workstations with the latest NVIDIA RTX PRO GPUs.

GPU acceleration sits underneath nearly every stage of reality capture, not just the AI-driven pieces getting most of the attention. Parallel processing on the GPU is what turns massive lidar point clouds and photogrammetric image sets into usable 3D data in a fraction of the time older sequential processing would take, and that same parallel horsepower drives the ray-traced rendering that makes scanned environments look convincingly real once they're done processing. It's also the foundation that makes newer techniques like neural radiance fields and Gaussian Splatting practical at all, since reconstructing photorealistic scenes from a handful of images is a massively parallel workload by nature. For teams working with large scans, the ceiling on scene size usually comes down to how much video memory the GPU has on board, not just how fast it runs, so more VRAM directly translates to bigger, more complex scenes a workstation can handle.

That is the broader message worth sitting with. AI experimentation in geospatial does not necessarily require a wholesale hardware investment before you can get started. Many teams are already closer to the starting line than they realize. The GPU specced for your current applications is also a parallel compute device, and the tools to put it to work on AI-assisted workflows are already available.

The question is not whether to engage with these technologies. It is how quickly the hardware can be matched to the ambition of the work.

Operational Benefits of Local AI

Centralized cloud-based AI compute operates on a consumption model. Every inference, every API call, every token processed carries a cost. For occasional use or small-scale experimentation that can seem manageable, but as AI becomes more deeply embedded in daily workflows those costs accumulate quickly and sometimes unpredictably. A team that starts using AI-assisted point cloud segmentation or automated feature detection at any real volume will find that the economics of centralized compute scale in ways that are difficult to forecast and harder to control.

Beyond cost, there is the question of operational risk. Geospatial professionals frequently work in environments where reliable connectivity cannot be taken for granted. A cloud-dependent AI workflow that encounters a network outage on a live job site is not just an inconvenience. Depending on the application, it can mean lost time, compromised data, or decisions made without the insight the workflow was supposed to provide. Local compute eliminates that dependency. The processing happens on the machine in front of you, regardless of what the network is doing.

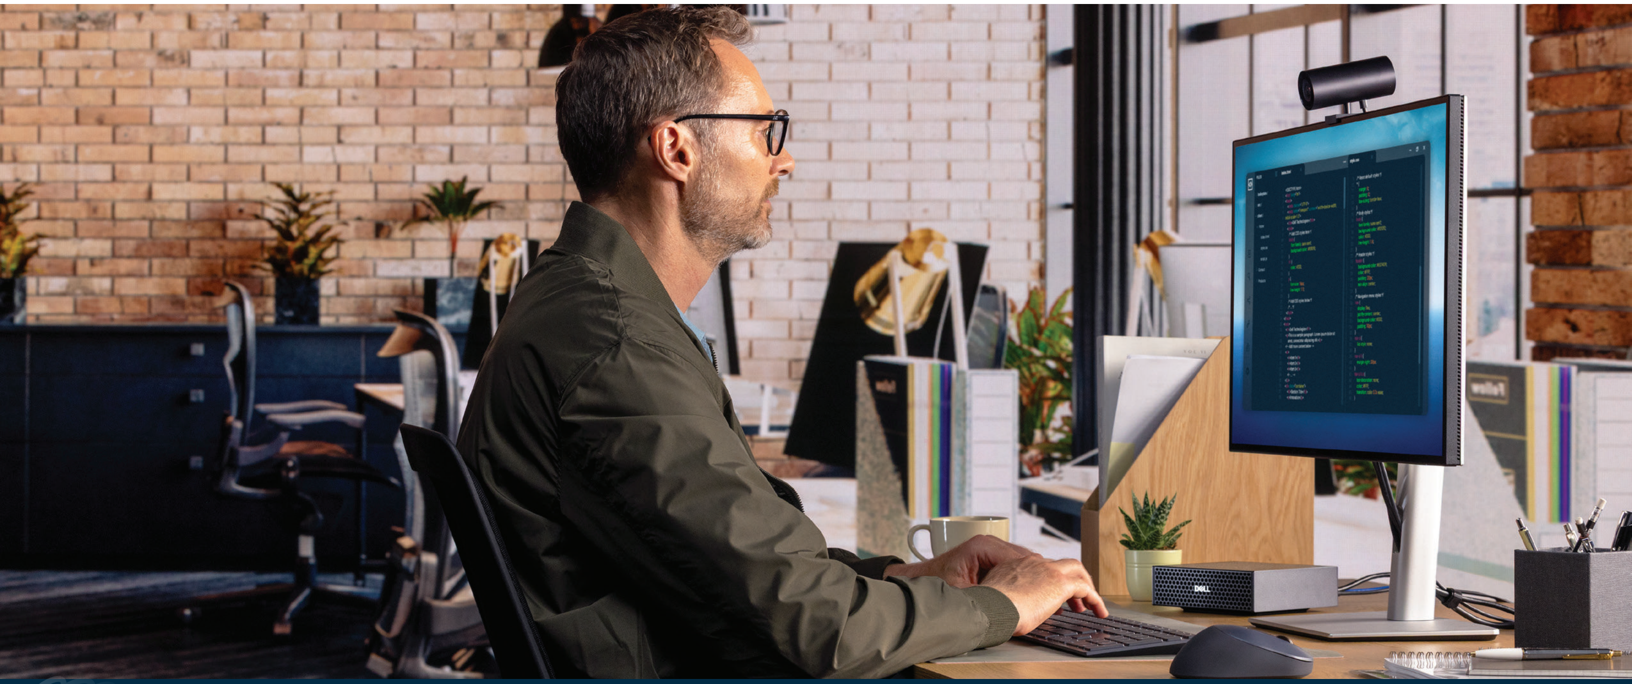

Running AI locally also changes the economics of getting started. There are no API costs to model, no consumption budgets to forecast, and no minimum commitment before a team can begin experimenting meaningfully. As agentic AI tools become part of professional workflows, that local foundation takes on even greater strategic value. Local accelerated compute can act as an AI offload engine for task-specific agents, classifying site data, preparing project summaries, flagging anomalies, generating design alternatives and coordinating multi-step analysis across local applications and datasets, all without a gloud dependency or a per-inference cost. The workstation becomes the foundation, and the investment in hardware is an investment in capability that compounds over time rather than a recurring operational expense that grows with usage.

For organizations weighing how to build AI into their geospatial workflows, that distinction matters. A solid local foundation does not just reduce risk and control costs. It creates the conditions for confident, sustained adoption, starting with what the team already has and scaling deliberately from there.

Product Spotlight:

The New AI Supercomputer

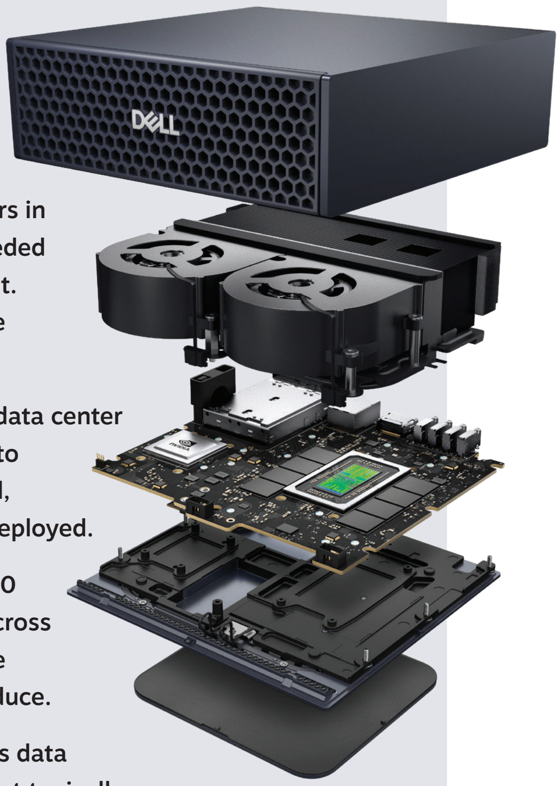

The Dell Pro Max with GB10 is a compact desktop device that, under the hood, is something closer to an AI supercomputer, purpose-built around the NVIDIA GB10 Grace Blackwell Superchip and running NVIDIA DGX OS. Originally designed with AI developers in mind, it was built for those training and deploying models who needed a desktop that could scale seamlessly to a data center environment. The geospatial industry is now discovering that the same hardware serves their workflows exceptionally well.

Because the GB10 shares its Blackwell architecture with NVIDIA's data center hardware, the workflows and models a team develops on the deskto today are fully compatible with larger infrastructure down the road, without requiring a fundamental revamp of how the AI is built or deployed.

The memory story is just as significant. The Dell Pro Max with GB10 features 128GB of LPDDR5X unified memory, coherently shared across both the CPU and GPU via NVLink-C2C technology, eliminating the host-GPU bottlenecks that traditional memory architectures introduce.

For AI workloads and large-scale geospatial processing, that means data moves freely between compute resources without the overhead that typically slows things down. In practical terms, the system can load models of up to 200 billion parameters entirely into memory and deliver up to one petaflop of FP4 compute performance, bringing a class of AI capability to the desktop that until recently lived exclusively in data centers.

For geospatial teams looking for a serious local AI capability without the footprint of a traditional high-end workstation, the Dell Pro Max with GB10 represents a genuinely new category of device. Compact, powerful, and built for the direction the industry is heading.

Local AI, Cloud AI, and the Case for Hybrid

Local and cloud compute are not competing philosophies. For most organizations, the smartest AI strategy draws on both. Cloud infrastructure remains the right choice for large-scale model training, elastic workloads, and collaborative workflows where data needs to be accessible across distributed teams.

Local compute earns its place where latency, connectivity, data volume, and cost predictability matter most, which for geospatial professionals working in the field describes most of the working day. The emerging best practice is a hybrid model, with local compute handling the time-sensitive and data-intensive parts of the workflow, and cloud infrastructure available for training, storage, and collaboration when needed. Dell Pro Max and Precision workstations with the latest NVIDIA RTX PRO GPUs, including the Dell Pro Max with GB10, are built to fit naturally into that architecture, delivering serious local AI capability while remaining fully compatible with cloud and data center infrastructure as needs evolve.

Conclusion: Future-Proofing Workflows

The geospatial industry has always been defined by the quality of its data and the speed at which that data can be turned into something actionable. What is changing now is the scale of what is possible, and the role that local compute plays in getting there. AI is not a feature on the horizon for this industry. It is already embedded in the tools teams use every day, and its footprint is only going to grow.

The workstation has come back to the center of that conversation for good reason. As AI workflows become more demanding, more data-intensive, and more time-sensitive, the case for processing power that lives close to the work, rather than somewhere out in the cloud, becomes harder to ignore. A survey team that can process lidar and point cloud data locally before leaving the job site leaves with insights, not just raw data. An AEC team running local AI to identify objects, defects, or safety issues from collected data can get answers in the field rather than waiting for a cloud pipeline to complete. A geospatial analyst running segmentation and classification locally can avoid the time, cost, and connectivity demands of uploading massive datasets to remote services entirely.

The teams that build a solid local compute foundation now will be the ones who can adopt new capabilities quickly, scale their workflows confidently, and avoid the operational and financial risks that come with depending on centralized infrastructure for mission-critical processing.

That foundation looks different for different teams. For some, it starts with the hardware already on the desk, a capable NVIDIA GPU that may already be ready for AI-assisted workflows with the right software and configuration. For others, it means stepping up to Dell Pro Max and Precision workstations with the latest NVIDIA RTX PRO GPUs, purpose-built for the accelerated compute demands of modern geospatial work. And for teams ready to push the boundary of what a desktop device can do, the Dell Pro Max with GB10 Grace Blackwell Superchip brings data center-class AI capability to a form factor that can go anywhere the work does.

The direction the industry is heading is clear. The workstations built to take it there are ready.

The question for every geospatial team is simply where they want to start.