Produced by
AECMAGAZINE

REALITY MODELLING SPECIAL REPORT

Reality modelling is unlocking new efficiencies in AEC projects, providing highly detailed 3D models across the entire lifecycle. To keep projects on track, rapid data processing is critical, so selecting the right workstation hardware is more important than ever

ALCATRAZ REALITY MODEL COURTESY OF PETE KELSEY, VCTO LABS

PHYSICAL MEETS DIGITAL

Reality modelling is reshaping the AEC sector, providing precise 3D models that bridge the physical and digital realms, improving planning, design, construction, and maintenance of projects

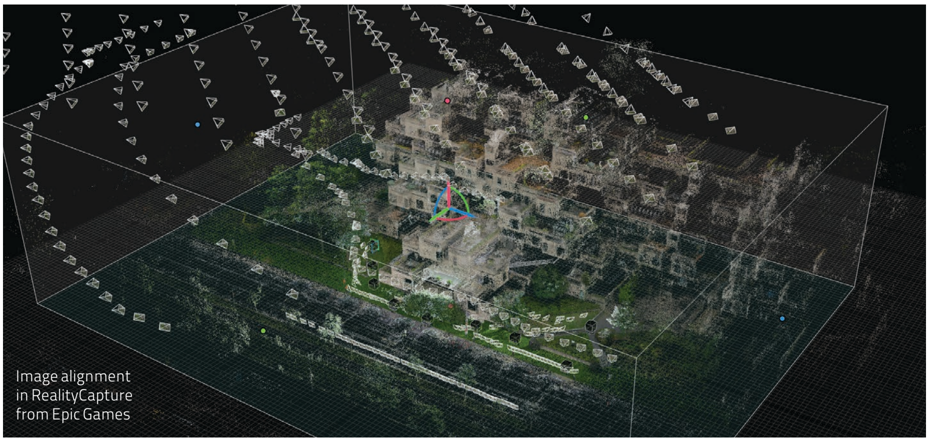

Image alignment in RealityCapture from Epic Games

Reality modelling is transforming architecture, engineering, construction (AEC), delivering highly accurate, data-rich 3D models from planning and design, all the way through to maintenance and operations. It's being adopted at an unprecedented rate and is set to become a cornerstone of modern construction, bridging the physical and digital worlds like never before.

The applications are incredibly diverse. Reality modelling can provide critical context for new projects, where accurate 3D site data allows designers to integrate new buildings or infrastructure into existing environments more effectively. It also serves as a foundation for retrofit projects, enabling precise planning for renovation and refurbishment.

Tracking construction progress has become more efficient as frequent reality capture helps teams compare current conditions against project schedules.

Construction verification is enhanced by comparing 'as-built' conditions with 'as-designed' BIM models, allowing discrepancies to be identified early, reducing costly errors. Additionally, reality modelling supports ongoing analysis of settlement, erosion, crack propagation, and other structural or environmental changes.

The technologies used to capture on-site reality are more advanced and accessible than ever before. Tripod-mounted terrestrial laser scanners, which capture precise 3D point clouds of buildings and infrastructure, are now augmented with photogrammetry and handheld SLAM scanners. These versatile scanners can also be mounted on backpacks, drones, or autonomous robots.

Multirotor drones are most common due to their stability and manoeuvrability, even in tight spaces, making them ideal for photogrammetry and close-range LiDAR scanning of buildings and infrastructure. Fixed-wing drones are better suited to large-scale mapping, such as surveying vast construction sites or road networks. Meanwhile, autonomous quadruped robots equipped with scanning technology can capture data on large or complex sites, even outside of standard working hours.

On a single project, reality capture can generate terabytes of raw data. One of the biggest challenges is processing this into usable data, such as registered point clouds, 3D meshes, or even BIM models.

All of these processes are computationally intensive, and the speed at which data is processed is critical to construction efficiency and decision-making. Fast processing ensures that stakeholders have up-to-date information, minimising delays and reducing costly re-work. In fast-paced construction environments, delays in converting reality capture data into usable models can lead to outdated site conditions, misaligned planning, and inefficiencies in project execution. Rapid processing allows for real-time monitoring of construction progress, enabling teams to detect and address deviations from design specifications before they escalate into major issues. It also enhances safety by identifying hazards or structural concerns as they arise, ensuring prompt intervention.

VISUALISATION - BRINGING REALITY MODELS TO LIFE

Reality modelling produces massive datasets, but raw data alone can be challenging to interpret. This is where visualisation plays a crucial role, transforming complex models into interactive real-time environments that enhance decision-making.

Tools like Twinmotion and Unreal Engine from Epic Games bring reality models to life with high-quality, real-time rendering. This greatly improves communication and makes it easier for stakeholders, clients and project teams to understand site conditions, identify issues early and make informed decisions faster.

Bringing reality models into a real-time viz environment also allows for seamless integration with design data from BIM authoring tools such as Revit, placing proposed buildings in their real-world context, even across cities.

Visualisation of London's skyline using Cesium's 3D geospatial technology. Image courtesy of Bentley Systems

REALITY MODEL BOTTLENECKS

The speed at which reality capture data is processed is critical to project efficiency. We highlight seven key photogrammetry, point cloud and mesh workflows where faster processing can make a huge difference

Point cloud registration in Leica Cyclone Register 360

1 IMAGE PRE-PROCESSING

In photogrammetry, captured images typically must be enhanced before they can be used to generate a point cloud or mesh. This often involves techniques such as tone stretching and denoising. To improve efficiency these can be automated and batch-processed using a workstation's CPU or (sometimes) GPU.

2 IMAGE ALIGNMENT

This is one of the most important steps in photogrammetry, where the software detects and matches key features across overlapping images, estimates camera positions and orientations, then creates a sparse point cloud. The process is computationally intensive and is often accelerated by multi-core CPUs.

3 FILE CONVERSION

Laser scan data often needs to be converted to different formats for interoperability, compliance, visualisation, or efficient storage through compression.

The process is computationally intensive but usually only runs on a single CPU core within a workstation, making processor frequency a priority.

4 POINT CLOUD REGISTRATION

Point cloud registration is a key workflow for laser scanning and involves aligning and merging multiple point clouds into a single, unified coordinate system.

The process is typically accelerated by a multi-core workstation CPU, though certain computational phases can be offloaded to the GPU.

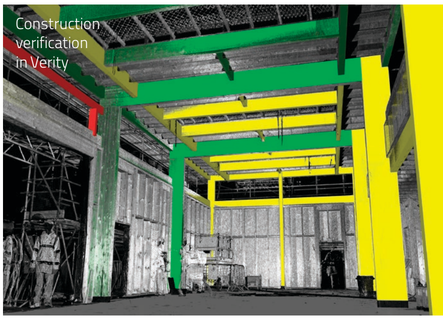

Construction verification in Verity

5 SCAN-TO-MESH

Compared to point clouds, mesh models are easier to interpret, making them ideal for visualisation. They also have smaller file sizes, which improves sharing and navigation. However, converting a point cloud into a mesh is computationally intensive. The process can be accelerated by multi-core CPUs and, in some cases, GPUs.

6 ANALYSIS

Various analyses can be performed on reality modelling data, such as comparing point clouds to monitor construction progress or structural changes over time, or aligning a point cloud with a BIM model for construction verification. These processes are typically handled by the CPU and are largely multi-threaded.

7 AI CLASSIFICATION

Auto-classification uses AI algorithms to intelligently categorise point clouds. The machine learning models are trained on large datasets of terrestrial scans to identify various elements such as floors, doors, ducts, pipes, beams, roads, and curbs. The process typically requires a GPU, but also heavily relies on the CPU.

Visualisation isn't limited to 2D displays — it can also be experienced in VR, allowing teams to immerse themselves in virtual spaces for planning, construction, and design review.

However, bringing reality models into real time viz tools comes with challenges. Maintaining high accuracy while ensuring smooth real-time performance is critical, especially for VR, where high frame rates are essential. Large, high-resolution meshes — often containing billions of polygons — can be difficult to render efficiently in real time, even with the most powerful workstation GPUs.

To optimise these datasets, models must be structured effectively. One approach is to simplify geometry, split models into multiple parts, and assign Levels of Detail (LoD), but this process can be time-consuming, restricting when visualisation can be effectively used throughout a project.

Another powerful option in Unreal Engine is Nanite, a virtualised mesh technology that dynamically streams and processes only the visible details, eliminating the need for manual LoD adjustments. However, converting a standard mesh into a Nanite mesh is computationally intensive. To streamline the scan-to-visualisation pipeline, a high-performance workstation with a powerful CPU is essential.

For large-scale environments, Cesium 3D Tiles provide an alternative solution, enabling massive 3D datasets to be streamed from a local resource or the cloud and rendered dynamically.

By using optimisation techniques, large reality models can be efficiently visualised in real time on powerful workstations, enabling smooth navigation and high-fidelity rendering — all without sacrificing performance.

CAPTURING THE ROCK

For Pete Kelsey of VCTO Labs, a project to capture Alcatraz in its entirety represented a once-in-a-lifetime opportunity. With a tight three-week schedule, processing data on-site was crucial — made possible only by the Lenovo ThinkStation P8 workstation and its 96-core AMD Ryzen™ Threadripper™ PRO processor.

There must have been some points during the long December nights spent on D Block at Alcatraz Island when Pete Kelsey questioned the series of decisions that had led him to occupy Cell 31.

But it's safe to assume that what kept him going was not just his infectious enthusiasm for reality capture technologies and projects, but also the once-in-a-lifetime opportunity to capture Alcatraz in its entirety and see the data put to work on preserving this historic landmark.

Kelsey has loved the National Parks of the United States since childhood and has worked pro bono on a number of important reality capture projects for the US National Park Service (NPS), including capturing the USS Arizona Monument at Pearl Harbor in Hawaii, the fossil wall at Dinosaur National Monument in Utah, and the USS Constitution in Boston.

"Alcatraz was a good fit for me," he says. "It's a great story, everybody's heard of it, and at 22.5 acres, the size is no issue for reality capture." So when he found out through conversations with NPS employees that they needed a way to start measuring the effects of sea-level rise on Alcatraz Island, he jumped at the chance to contribute. Reality capture, he argued, would create a baseline survey against which future studies of the Island could be compared, in order to quantify those effects over time.

Alcatraz closed as a prison in 1963, having become too expensive to maintain. It opened to the public in 1973 and is now a major tourist site that attracts some 1.2 million visitors per year.

Erosion has taken its toll on the Island and continues to do so. While it's a natural geological process, it has caused the Island's cliffs to recede and walkways to collapse into the San Francisco Bay. "This baseline survey could help NPS to triage repairs and maintenance work – basically, deciding what to fix first," says Kelsey. "And I knew from the start that it would necessitate scanning the whole island, right down to the waterline, preferably at low tide."

A PUNISHING SCHEDULE

This survey needed to be exhaustive in its level of detail. As Kelsey points out, once people know that the data from a reality capture project exists and is publicly owned, they tend to have a lot of questions to ask it. And those questions might lead to all sorts of valuable insights, depending on the inquirer. Geologists might wish to simulate the impact of a major earthquake on the Island. Marine biologists might want to analyse the nesting habits of seabirds there. "Once that data's captured and available, there's no putting the genie back in the bottle," Kelsey jokes.

At the same time, he and his team were facing a number of constraints regarding the project. Above all, they had only been granted a three-week window of time to perform the entire capture of Alcatraz. In any case, it was nine months before he could start the work with the necessary permits from the NPS in hand, because seabirds had to fledge and leave their nests before drones could be operated in the area.

A regular tide of visitors arriving and departing from the island was also a complication, even in December. It raised issues when it came to both flying drones and scanning interiors, and Kelsey was also mindful of the need to respect tourists' vacation time and not disrupt it.

Weather, too, was potentially a problem. Strong winds, choppy waters and dense fogs are year-round occurrences on the San Francisco Bay.

All this meant it was a priority for the project team to make the most of whatever access time was available to them. Kelsey quickly realised that meant that a core team of around eight people - including himself - would have to stay on the island, rather than rely on the boats that ferry employees and visitors across the Bay to the Island between early morning and 4.30pm. "The NPS told me, 'Sure, you can stay overnight if you want to – but you'll all have to sleep in cells.'"

Undeterred, Kelsey and his team pushed on with the project, using a vast array of reality-capture technologies. "There's no one piece of gear that could do everything we needed it to, to achieve the kind of comprehensive capture we were seeking. It doesn't exist," he says. "So that's why I had to create a sort of Frankenstein-esque collection of skills, people and technologies to get the work

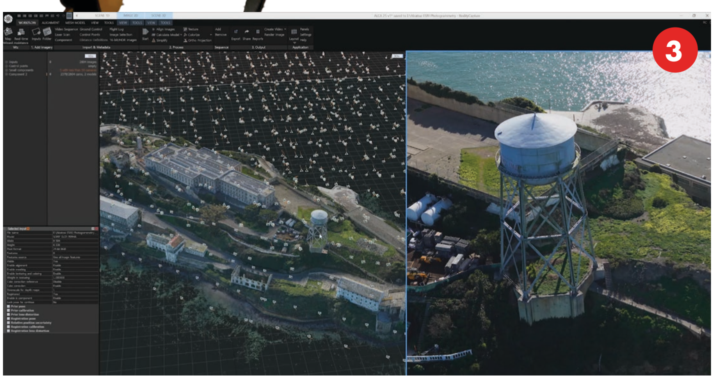

done in the three weeks available to us."

The first step was to set the survey control network - the markers or targets acting as fixed reference points on the island for sensors on drones. Kelsey had already decided on three types of drone-based capture technologies – LiDAR, photogrammetry and multi-spectral imagery – to achieve a complete view of the Island and building exteriors. With those plans in place, the two days set aside for actual flying could not have gone better, he says. "It was perfect, perfect weather – a miracle, really. And the work completed on those two days ticked off the biggest strategic aim for the project, in terms of capturing data to address sea-level concerns."

With the interiors, Kelsey settled on SLAM LiDAR technology. "The level-of-detail requirement meant terrestrial laser scanning wasn't necessary, and we'd probably still be there if we'd done it that way," he jokes.

"We were SLAM scanning from Day One, using handheld scanners you walk around holding. But they could also be hung from drones, which we did using an Elios 3 drone from Flyability, contained in a protective cage so it can bump into ceilings and walls without crashing. We mounted them to backpacks. We even used Spot, the mobile dog robot from Boston Dynamics, to take scanners into areas of the buildings that were totally off-limits, due to asbestos and lead contamination."

The team Kelsey used included representatives and equipment from a huge range of vendors, including Emesent, Riegl, Freefly Systems, Flyability and Boston Dynamics – all contributing their time and kit to the project on a pro bono basis.

POST-PROCESSING CHALLENGES

All post-processing was conducted in real time on the Island, where connectivity is a real challenge and the Wi-Fi network operates "at dial-up speeds", according to Kelsey. With terabytes of data to crunch through, this meant the cloud was simply not an option.

Team members would perform a scan, come to the cellhouse office and dump gigabytes of data for post-processing, there and then. "Before we left this remote location, we needed to know that we'd scanned everything we needed to, that the data captured was usable and that the post-processing had been successful," he says.

Data was pumped into RealityCapture from Epic Games. "I knew I wanted to use this, because it's one of the only products I know about that can integrate both LiDAR and photogrammetric data in a single model. With that combination – the geometry of LiDAR and the photorealism of photogrammetry – the resulting models are really second to none," says Kelsey.

Help with this hefty compute burden soon arrived in the form of the Lenovo ThinkStation P8, a powerful 96-core AMD Ryzen™ Threadripper™ PRO workstation with 256 GB of memory. The machine shouldered the burden immediately and, according to Kelsey, it literally saved the project.

"I will never forget that day in the office at Alcatraz, with RealityCapture crunching away on our capture data, probably the photogrammetry," Kelsey recalls. "I brought up the Task Manager to see 96 blue squares, all running at 100% at 4.7 GHz. I don't mind admitting that I squealed with excitement."

Soon, he was confidently running simultaneous workloads; for example, with SLAM scanning data captured using Emesent's Hovermap devices and Aura software: "We found we could get 12 to 15 sessions of Aura running on this monster of a workstation at the same time. This was a huge benefit to our schedule."

In fact, that workstation was a big factor in the project's successful, on-schedule finish. The result is a unique, 3TB dataset offering multiple layers of detail, thanks to the diverse tools used. It is capable of supporting all manner of applications, from efficient park management to historical research. And it doesn't stop there. The reality mesh and point cloud datasets are currently being combined in Unreal Engine to develop a hyperealistic, immersive real-time experience.

But the thing that Kelsey is perhaps most looking forward to is sharing his passion for Alcatraz Island with the rest of the world. "We now have the data for the NPS to create a virtual Alcatraz that anyone can visit, regardless of where they live. I can't say when or if that will happen, but the potential for outreach and education in the wider world is inspiring."

1 Pete Kelsey used the Lenovo ThinkStation P8 workstation with 96-cores for extreme multi-tasking

2 Spot, the mobile dog robot from Boston Dynamics, took the Emesent Hovermap LiDAR scanner into areas of the buildings that were off limits

3 RealityCapture from Epic Games was used to process the bulk of the data

4 Detailed LiDAR scan of the Alcatraz Island site, captured by drone

5 The AMD powered Lenovo ThinkPad P16s laptop was used for drone flight planning using Sitescan for ArcGIS

STORIES AND SECRETS

At HOK's Centre Block project in Ottawa, the race is on to capture 25,000 historical assets and the building in which they reside, with stakeholders able to explore the massive reality model via an Unreal Engine hub.

In the multi-year project to preserve, restore and modernise one of Canada's most iconic buildings, HOK director of design technology Mark Cichy has learnt to expect the unexpected.

After all, Centre Block is not just the main building of the Canadian parliamentary complex on Parliament Hill in Ottawa, Canada. It's also a physical manifestation of the nation's democratic history, with many stories – and secrets – to tell. Every day brings new surprises, says Cichy: antique mouldings get uncovered, hazardous materials are found, animal remains get dug up.

"But the important thing is that there should be no surprises for the client," he says. From the start, the project has involved painstaking work to capture not only the building and its spaces, but also more than 25,000 movable and fixed heritage assets found inside, ranging from radiators to works of art.

"On this project, there is significant investment in recording past states," Cichy explains. "So, a key client requirement is that we compile a complete record of the building over time, capturing each and every space before anyone moves in to dismantle it, and then while it is being dismantled, and then while rooms and assets are being remediated and reinstalled."


Aerial view of Canadian parliamentary complex


Unreal Engine visualisation

This work calls for a wide range of reality capture technologies. LiDAR drones are used for capturing views of the site and of Centre Block's exterior. Inside the building, tripod-mounted laser scanners are used, supplemented by SLAM devices in tight spaces or areas that are otherwise difficult to access.

Reality capture data, in the form of both photogrammetry meshes and point clouds, are combined with architectural, structural, and MEP design content from Revit and Rhino.

Due to Centre Block security restrictions, it's not permitted to use wireless networks or personal devices within the building itself, so reality capture data must leave the building on security-checked capture devices and then be physically transported to the project team's integrated project office a short distance away from the main site, where the data is processed. AI plays a big part in streamlining data processing, identifying hazards, and automating documentation.

Data is then pumped into a massive hub powered by Unreal Engine, which is capable of ingesting many terabytes of data in real time. For interactive visualisation, Nanite technology dynamically renders only the visible details, eliminating the need for manual Level of Detail (LoD) adjustments. Everything is processed on the hub, with pixels streamed to any device with an internet browser.

The result is a dynamic resource that enables multiple project stakeholders to concurrently explore highly detailed 3D models at various stages, assess the impact of design choices in real time, take VR walkthroughs and interactively review and provide feedback on the project's progress.

The computational demands are considerable. To test and scale the Unreal Engine build, HOK is currently using a multi-GPU Lenovo ThinkStation workstation as a "sandbox" environment.

It's an impressive achievement, both in terms of technology deployment and historic preservation. And Cichy sees the Unreal Engine-based hub continuing to play a valuable long-term role in how the Canadian government uses Centre Block, long after the project is completed, in areas such as facility management, building operations, office planning and staff distribution.

CRUSHING THE COMPUTATION

Reality modelling is one of the most computationally demanding tasks in AEC. With the right workstation hardware, teams can save hours, if not days, of processing time, accelerating project timelines.
We break down the key components and highlight what to consider when specifying a workstation.

GPU (GRAPHICS PROCESSING UNIT)

In reality modelling, the Graphics Processing Unit (GPU) serves two key roles: 3D model visualisation and data processing. Visualisation can take place within dedicated reality modelling applications or in real-time game engine environments like Unreal Engine, where the GPU faces greater demands due to the emphasis on realism and interactivity, especially in VR.

For data processing tasks such as mesh reconstruction or AI classification, more powerful GPUs can speed up workflows. However, because these processes share some of the workload with the CPU, the performance gains from a high-end GPU are often less significant compared to a fully GPU-driven task like ray-traced rendering. As a result, entry-level workstation GPUs can often provide better value for money than pricier high-end models.

For reality modelling there is currently little reason to invest in more than one GPU unless, perhaps, the workstation is virtualised to serve multiple users.

STORAGE (SSD)

Fast NVMe Solid State Drives (SSDs) are crucial for handling massive reality modelling datasets, often reaching terabytes in size. Processing this data can also generate huge temporary files, making high-endurance SSDs even more essential.

To optimise performance, separate SSDs are often recommended — one for reading, one for writing, and sometimes a third for the operating system and applications. Advanced tower workstations, such as the Lenovo ThinkStation P8, support multiple SSDs directly on the motherboard and also offer front-accessible SSDs and PCIe add-in boards that can host multiple drives.

For additional flexibility, multiple SSDs can be configured in RAID arrays. RAID 0 (striping) enhances performance, while RAID 1 (mirroring) provides redundancy, safeguarding data in case of a drive failure. Given that it can take days to process the most complex reality modelling datasets, this serves as a shrewd insurance policy. Meanwhile, traditional Hard Disk Drives (HDDs) remain the most cost-effective option per gigabyte, making them ideal for archiving.

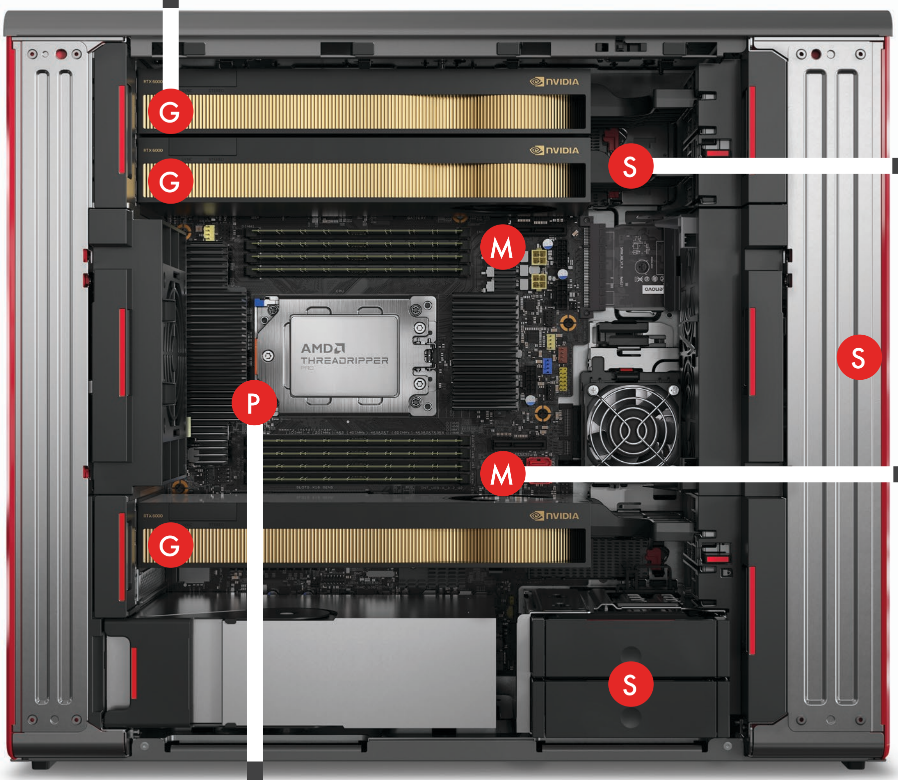

Lenovo ThinkStation P8 with AMD Ryzen Threadripper PRO processor

MEMORY (RAM)

Workstation memory can have a big impact on performance in reality modelling software. When handling large datasets, or running multiple operations in parallel, running out of memory can significantly slow performance, forcing the workstation to temporarily page data to the SSD to complete calculations.

Memory performance is equally important, especially when multi-tasking. It is determined by memory bandwidth, which depends on both the number of supported memory channels, and the frequency of the memory modules. AMD Ryzen™ Threadripper™ PRO processors offer a strong balance of both, supporting 8-channel memory with speeds of 4,800 MHz. To fully utilise this bandwidth, the number of memory modules should match the number of available channels.

Additionally, ECC memory, standard on AMD Ryzen Threadripper PRO, helps prevent crashes — critical for lengthy calculations where losing hours, or even days, of newly processed data is not an option.

PROCESSOR (CPU)

Many reality modelling workflows are multi-threaded, so benefit from multiple CPU cores. However, more cores can mean diminishing returns, especially since certain stages in some workflows rely on just one core. Other tasks are entirely single-threaded, making it essential to strike the right balance between core count and clock speed (GHz).

CPUs with large caches — high-speed on-chip memory for frequently accessed data — can further enhance performance. Additionally, AMD Simultaneous Multithreading (SMT), which enables a single core to run multiple threads, can impact processing time. Disabling it can sometimes lead to faster execution.

AMD Ryzen Threadripper PRO processors offer a powerful combination of high core counts, fast clock speeds, and large caches. With models ranging from 12 to 96 cores, they cater to a variety of workloads. When selecting a processor, multi-tasking should also be considered, as processing data in parallel from multiple sources in multiple applications, can provide significant productivity gains.

REALITY TECH: WORKSTATIONS

Lenovo workstations powered by AMD processors, including the AMD Ryzen™ Threadripper™ PRO with up to 96 high-performance cores, make light work of the most demanding reality modelling workflows

Reality modelling presents some of the most demanding workflows in AEC. And as reality capture hardware continues to improve in resolution, datasets grow larger, and capture frequency increases, these demands continue to intensify.

Processing LiDAR or photogrammetry data can take hours, even days. On a busy construction site, every second counts, making fast, accurate data delivery essential for decision-making, minimising delays, ensuring quality control, and reducing costs.

Powered by the AMD Ryzen™ Threadripper™ PRO processor, the Lenovo ThinkStation P8 is engineered to handle the most demanding reality modelling workloads. It delivers the speed and efficiency needed to keep projects on track, whether in the office or on-site.

What sets the ThinkStation P8 apart is its versatility — offering a wide range of AMD Ryzen Threadripper PRO processors, that simplify fleet management for IT admins. Choose from 12, 16, 24, or 32 cores for mainstream workflows, or scale up to 64 or 96 cores for the most demanding workloads and extreme multi-tasking.

With up to 1TB of DDR5 memory, the P8 seamlessly handles massive datasets, while up to 8 M.2 SSDs provide extensive storage options, allowing users to prioritise both performance and redundancy.

With up to 10Gbe Ethernet as standard and optional 25Gbe, it also enables superfast transfer of colossal reality modelling datasets across the network.

Reality modelling doesn't have to be confined to the office. The P8 can be rack-mounted for remote access, delivering huge computational power from any location.

For those who need powerful computational capabilities on the go, Lenovo ThinkPad mobile workstations are an excellent choice. The ThinkPad P14s and ThinkPad P16s strike the perfect balance between performance and portability. While they may not match the scalability of the ThinkStation P8, they still pack a serious punch with AMD Ryzen 7 PRO 7040U Series processors offering up to 8 cores and 64GB of RAM.

Lenovo workstations are more than just high-performance hardware — they are built to withstand the rigours of professional use. With legendary build quality, superior thermal management, and impressive reliability, these systems are designed to run cool, quiet, and consistently, even in the most demanding environments.




Lenovo ThinkStation P8	Lenovo ThinkStation P620	Lenovo ThinkPad P14s Gen 4 AMD	Lenovo ThinkPad P16s Gen 2 AMD
• Flexible and expandable high-end desktop or rack-mounted workstation for multi-application, multi-threaded reality modelling workflows with huge datasets • Can be optimised for CPU or GPU accelerated workflows, including high-end visualisation	• Value-driven flexible and expandable high-end desktop workstation for multi-application, multi-threaded reality modelling workflows with huge datasets • Can be optimised for CPU or GPU accelerated workflows, including high-end visualisation	• Compact, lightweight, sturdy, and energy efficient 14-inch (2.8K) mobile workstation for entry-level reality modelling workflows, especially on construction sites • Emphasis on CPU computational workflows, with GPU used for viewing reality models	• Compact, lightweight, sturdy, and energy efficient 16-inch (4K) mobile workstation for entry-level reality modelling workflows, especially on construction sites • Emphasis on CPU computational workflows, with GPU used for viewing reality models
AMD Ryzen™ Threadripper PRO 7000 WX Series processor with 12, 16, 24, 32, 64 or 96 cores	AMD Ryzen™ Threadripper PRO 5000 or 3000 WX Series processor with 12, 16, 24, 32, or 64 cores	AMD Ryzen™ 5 / 7 PRO 7040U Series processor with 6 or 8 cores	AMD Ryzen™ 5 / 7 PRO 7040U Series processor with 6 or 8 cores
Up to 1 TB (8 x 128 GB) DDR5 4800 3DS-RDIMM ECC memory 3DS-RDIMM) memory	Up to 1 TB (8 x 128 GB) DDR4 3200 RDIMM ECC memory	Up to 64 GB LPDDR5 memory at 6,400 MT/s speeds	Up to 64 GB LPDDR5 memory at 6,400 MT/s speeds
Up to eight drives (8 x M.2 SSD) or (4 x M.2 SSD + 3 x 3.5" SATA HDD)	Up to 11 x M.2 SSD + 5 x 3.5" SATA HDD	One drive, up to 2TB M.2 SSD	One drive, up to 2TB M.2 SSD
Up to AMD Radeon™ PRO W7900 or NVIDIA RTX 6000 Ada Generation GPU	Up to AMD Radeon™ PRO W7900 or NVIDIA RTX 5000 Ada Generation GPU	Integrated AMD Radeon™ 740M or AMD Radeon™ 780M GPU with shared memory	Integrated AMD Radeon™ 740M or AMD Radeon™ 780M GPU with shared memory

Learn more
The Lenovo workstation portfolio
www.lenovo.com/amd-workstations