

# Workstation virtualisation

## *AMD vs Intel vs Nvidia*

With its new 'POC in a week' offering, IMSCAD is providing firms that use CAD and BIM tools with a fast-track virtualisation trial and a chance to see which GPU technology works best for them.

by GREG CORKE

On-site data centre virtualisation deployments capable of streaming 3D CAD to thin clients are notoriously complex. With every firm using a different mix of applications, datasets, workflows and IT infrastructure, one-size-fits-all solutions are rare. Proof of Concepts (POCs) often drag on for months. Some never even get off the ground, with server hardware left in cardboard boxes gathering dust.

That's the experience of Adam Jull, CEO of graphics virtualisation specialist IMSCAD, which offers design, engineering and architecture firms a much easier way to see if virtualisation can work for them.

The UK-based consultancy's new 'POC in a week' offering is designed to take the pain out of getting up and running with hardware, software and licensing.

Rather than an on-premise install, firms instead get remote access to a bespoke, off-site data centre solution that supports five users on up to five CAD applications over a five-day period.

IMSCAD advises clients on the type of deployment – whether they should virtualise specific applications, for example, or deliver a fully virtual desktop to users. Solutions can be built around Citrix or VMware, Windows 7/8 or Windows 10.

Importantly, and (we think) uniquely, IMSCAD allows firms to try out all three graphics virtualisation technologies capable of delivering graphical accelerated applications, including 3D CAD.

As IMSCAD does not sell graphics or server hardware, Jull says the company is in a position to give firms independent advice on which GPU technologies – AMD MxGPU, Intel Iris Pro or Nvidia GRID – are best suited to their workflows, applications, datasets and user requirements. After an initial phone consultation, it will then configure one or more bespoke POC servers, complete with software. IMSCAD charges a \$5,000 fee for this service to cover administration and consultative support.

The company has servers in London and the USA (both West Coast and East Coast), so it can serve all of North America and Europe. It also has a pool of software licences, including Siemens NX, Solid Edge, SolidWorks, Catia, Revit, AutoCAD, Inventor, Civil 3D and other Autodesk software. Other applications can be delivered and customers can install their own as well.

Customers upload their CAD files to the data centre and then access the 'workstations' from software clients on PCs, laptops or tablets, at home or in the office, over a standard Internet connection.

If bandwidth is poor or latency is high, WAN acceleration technologies like Citrix Cloudbridge or Riverbed can be used.

To help users get the most out of the week, IMSCAD does some hand-holding along the way and collates empirical feedback and statistical information.

User experience is assessed through response sheets, where CAD users give

marks out of ten for specific actions such as key modelling tasks and model manipulation (for example, rotate, zoom and pan).

Performance of the environment is constantly monitored using Goliath software, to help identify bottlenecks in the system and give IMSCAD important information as to how the POC might be scaled up to support more users.

At the end of the week, IMSCAD goes through the results and gives advice on what solutions could potentially work for each firm. Of course, it can then roll out a larger-scale deployment on its behalf, either on site or in the cloud.

### Graphics flexibility



Nvidia may be the established leader in graphics virtualisation with Nvidia GRID, but Jull believes that both Intel and AMD can now offer some healthy competition.

Intel's 'Skylake' Xeon E3-1500 v5 with Iris Pro graphics P580 is a quad-core CPU and GPU in one package. This means that there is no need for an add-in PCIe Nvidia or AMD GPU. In the data centre, a server typically has rows of these processors, each housed on a cartridge-based chassis, which is essentially a collection of self-contained mini workstations.

Every CAD user connects directly to their own workstation, so there is no need to get into the realms of graphics virtualisation. This, says Jull, makes configuration relatively easy.

“  
**Nvidia is a great solution but it's still costly and it's fairly complex to deploy. Any of the three solutions pose their own complexity, but as long as the end-to-end infrastructure is right, your solution will deliver**  
 ”

## IMSCAD - GPUs at a glance



GPU technology	AMD MxGPU	Intel Iris Pro graphics	Nvidia GRID
<b>Typical CAD applications</b>	AutoCAD, Revit, SolidWorks, Siemens NX	AutoCAD, Revit	AutoCAD, Revit, SolidWorks, Siemens NX, 3ds Max, Nvidia Iray
<b>GPU models</b>	AMD FirePro S7150 or FirePro S7150x,	Intel's 'Skylake' Xeon E3-1500 v5 with Iris Pro graphics P580	Nvidia Tesla M60
<b>Form Factor</b>	PCIe card (single or dual slot)	GPU embedded inside CPU	PCIe card (dual slot)
<b>GPU virtualisation</b>	Hardware	N/A	Software
<b>Virtualisation software support</b>	VMware vSphere/ESXi 6.x or later, VMware Horizon 7.x or later (Citrix support coming in 2017)	Citrix XenServer 7.0 or later or XenDesktop/ App or Horizon/View	Citrix XenServer v6.5 SP1 or later, Citrix XenDesktop v7.8 or later, VMware vSphere 6 or later, VMware Horizon 7 or later
<b>Typical partner CPUs</b>	Dual Intel Xeon E5-2600 v5	N/A	Dual Intel Xeon E5-2600 v5
<b>Typical density for Revit (entry-level 3D CAD)</b>	16 users per 2U server	45x users per 4U server (1 user per CPU/ GPU) (HP Moonshot, Super Micro)	32 users per 2U server
<b>Typical density for SolidWorks (mid-range 3D CAD)</b>	8 users per 2U server	as above	16 users per 2U server
<b>Typical density for 3ds Max (high-end 3D)</b>	4 users per 2U server	not recommended	4-8 users per 2U server
<b>End User Licenses</b>	N/A	N/A	Yes
<b>Certified ISV Applications</b>	Yes	Yes	Yes

“With Intel, you don’t need a hypervisor. It’s a direct bare-metal delivery with direct XenDesktop/App or Horizon/View,” he explains.

It is possible to ‘virtualise’ Intel Iris Pro graphics, breaking up the GPU into multiple instances. However, this isn’t something he recommends for most graphical-based applications.

IMSCAD considers the ‘Skylake’ Intel Xeon E3-1500 v5 to be a significant step up from the ‘Broadwell’ Intel Xeon E3-1200 v4 with Iris Pro Graphics P6300 that it used in several POCs last year.

Jull says the ‘Broadwell’ Xeon was a little underpowered for most graphical applications, and that it lost out to Nvidia GRID in several POCs. “Each time it was close,” he says, “but Nvidia GRID won – graphics [performance] was the key issue for the actual users.”

With the promise of up to 30% more GPU power, Jull believes ‘Skylake’ Xeon will be a good fit for CAD and BIM. “I think Revit/AutoCAD workflows can be achieved on the Intel solution. Maybe not Siemens NX and SolidWorks and those sorts of [graphics-intense] applications – but certainly the architectural CAD market,” he says.

AMD made its long-awaited entrance to graphics virtualisation in September 2015 with the AMD Multiuser GPU (MxGPU), pitched as the world’s first hardware-based virtualised GPU solution.

Simply speaking, it means that the GPU is built from the ground up for virtualisation, with all the virtual machine assignments performed inside the silicon, rather than inside software, as is the case with Nvidia GRID.

AMD MxGPU uses SR-IOV (Single Root I/O Virtualisation) technology, a standard way for devices such as network adapters to expose hardware virtualisation.

AMD offers two PCIe add-in boards, the

single-height FirePro S7150 (one GPU) and dual-height FirePro S7150x2 (two GPUs). “For an AutoCAD-type workflow,” Jull says, “AMD can potentially support up to 32x graphical (CAD) VDIs on a server. If you’re doing SolidWorks/Catia it could be down to about eight.” This is based on there being two double-width GPUs inside a single server.

Jull explains that, because AMD’s GPUs are hardware virtualised, there is only one driver to apply and no software licensing [Nvidia charges a per-user, per-year licence for GRID]. This, he says, makes it slightly easier to deploy and brings down the total cost of ownership (TCO).

As users have access to a dedicated hardware resource, performance is also more predictable. “With Nvidia GRID, it can be more variable because you are creating virtual profiles on the GPU,” he says.

The downside of the AMD solution, according to Jull, is it is not as scalable and powerful as Nvidia GRID, depending on the user mix. Because AMD MxGPU uses hardware virtualisation, GPU resources are pre-sectioned off and you can’t go any finer than that.

“With Nvidia GRID – say, in production after three months – if you find there is spare capacity on the GPU, you could add another 16 users by choosing a different profile.”

Nvidia GRID scalability improved significantly last year with the introduction of the Tesla M60, a GPU accelerator that boasts twice the performance of the previous generation Nvidia GRID K2 GPU.

With a range of virtual GPU (vGPU) profiles, the Tesla M60 has the flexibility to support the widest range of users, from low end to very high end, and the highest densities. With a typical Nvidia GRID server with two GPUs, you could support around 32x Revit users or 16x SolidWorks users, Jull reckons.

“The options with Nvidia are vast. With [the company] adding the new M10 and M6 for blade servers, they really cover everything from Windows 10 to Catia,” he says, adding that the Tesla M60 can even be used for high-end compute, such as the physically based renderer, Nvidia Iray.

“Nvidia is a great solution but it’s still costly and fairly complex to deploy. Any of the three solutions pose their own complexities, but as long as the end-to-end infrastructure is right, your solution will deliver.”

“In the AutoCAD and Revit space – the architectural space – I think there is definitely a potential shift towards an AMD and Intel mix of solution,” he adds. “Ultimately, you could use all these technologies for your users at some level. This just provides a richer and varied choice of how to centralise your data.”

## Conclusion

For a typical POC, bringing technologies together for an on-site install can be a lengthy process – and one that is usually tied to a specific GPU vendor. With this initiative, IMSCAD is trying to streamline the process, to give users fast and valuable experience of working with their own CAD/design datasets, using one or more of the three key GPU technologies.

This is all about firms proving the technology can work for their users. Then, at the successful conclusion, IMSCAD can help them understand what is required to roll out a production environment.

At approximately \$5,000 per POC, firms need to be serious about investing in workstation virtualisation, but it’s a small insurance payment that reduces the risks associated with what is likely to be a much larger one.

■ [imscadglobal.com](http://imscadglobal.com).

## What is workstation virtualisation?

Read our beginner’s guide to workstation virtualisation in the September/October 2015 edition of AEC, or at: [tinyurl.com/WS-virtual](http://tinyurl.com/WS-virtual)

