

SUCCESS STORY | AWL-TECHNIEK

AWL-TECHNIEK EXPANDS DESKTOP VIRTUALIZATION GLOBALLY USING NVIDIA QUADRO vDWS



Image courtesy of AWL-Techniek



AWL LEVERAGES NVIDIA QUADRO vDWS FOR GRAPHICS-ACCELERATED COLLABORATION AROUND THE WORLD.

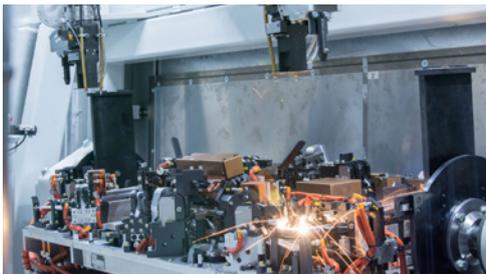


Image courtesy of AWL-Techniek

SUMMARY

- > AWL-Techniek is a world leader in the development and manufacturing of high-tech welding machines for the automotive and machine industries.
- > To improve productivity, AWL needed a high-performance virtual desktop infrastructure (VDI) environment that enabled engineers to collaborate on 3D graphics files.
- > Global expansion also meant AWL needed to control IT costs as staff numbers grew rapidly.
- > The company deployed NVIDIA® Quadro® Virtual Data Center Workstation (Quadro vDWS) software to deliver 3D graphics applications to engineers and improve productivity applications for office staff.

INTRODUCTION:

AWL-Techniek is a world leader in the development and manufacturing of high-tech welding machines. The company's engineers specialize in designing innovative concepts for high-end machine automation, robotization, and joining technologies for the unique requirements of international automotive suppliers. Thanks to AWL's stellar reputation, customer orders have skyrocketed in recent years as the automotive industry's global demand has increased.

CHALLENGE STATEMENT:

In response to growing sales, five years ago, AWL began expanding from a regional player into a global partner for its automotive customers by increasing its production capacity around the world. In 2013 the company had two facilities in the Netherlands and the Czech Republic. As it planned for new locations in China, the United States, and Mexico, AWL's leaders needed to find ways of maximizing efficiency and improving global collaboration across all business sectors. Overcoming several big IT hurdles was a key strategy for achieving these goals.

CUSTOMER PROFILE



Organization
AWL-Techniek

Industry
Manufacturing

Location
Harderwijk,
Netherlands

Employees
700+ employees

Website
www.awl.nl

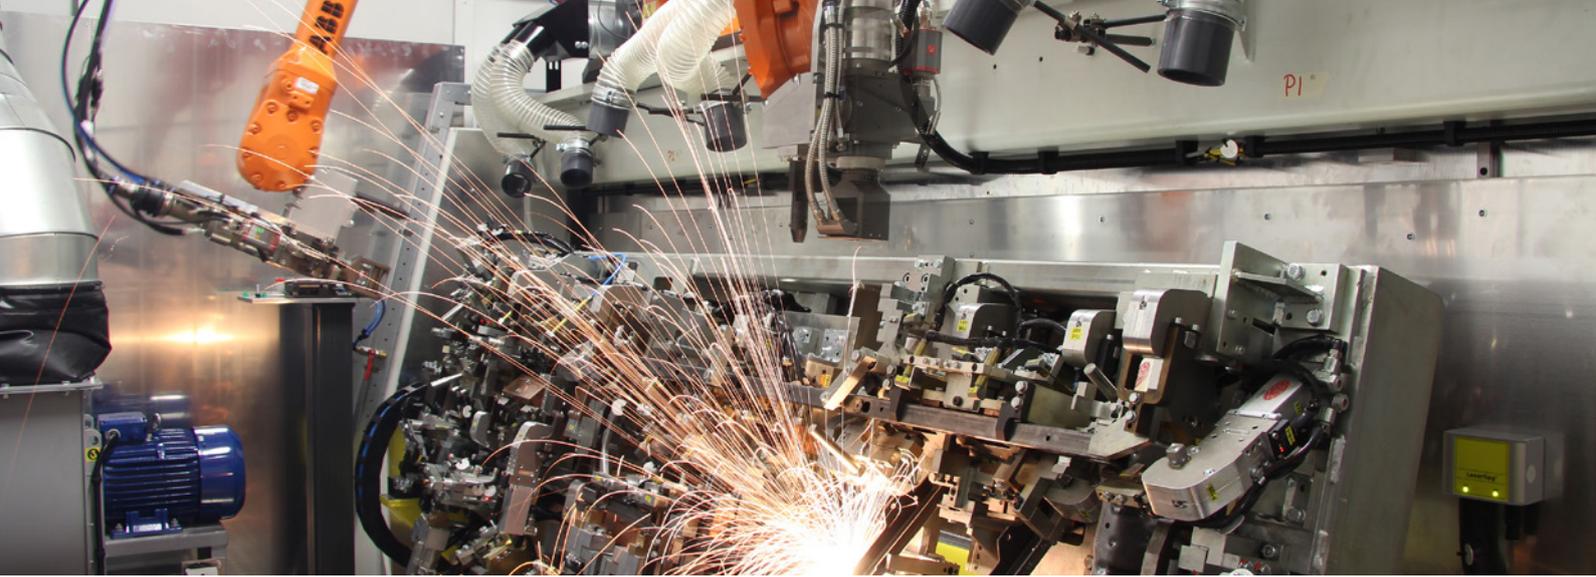


Image courtesy of AWL-Techniek

PRODUCTS

Hypervisor: Citrix XenDesktop, XenApp, and XenServer

Graphics Acceleration: NVIDIA Quadro vDWS

Server: Dell PowerEdge R720, R730, and R740

GPU: NVIDIA GRID® K2, NVIDIA Tesla® M60, and NVIDIA Tesla P40

RESULTS WITH NVIDIA QUADRO vDWS

- > Expanded global locations with 700+ staff with an IT team of 10 in the Netherlands and no local IT
- > Increased server density by 300 percent with users experiencing the same, or better, performance
- > Improved user experience with reduced end-to-end latency when leveraging NVENC
- > Decreased CPU usage by 28 percent with NVENC

At the time, one of AWL's biggest IT challenges was a lack of data synchronization. The company's mechanical design engineers faced daily productivity setbacks when working on shared 3D files in applications like Siemens PLM NX, Siemens Process Simulate, Dassault Systèmes CATIA, and SOLIDWORKS. To update files, engineers downloaded them from the server and uploaded them on completion. This process meant that engineers were never sure whether they were working on the latest revision of files.

When a project was ready for production, engineers exported the 3D files and transferred them to mechanics for viewing in eDrawings. This process introduced challenges with large data transfers and errors. It was common for engineers to update files and then forget to export a new viewable version or inform colleagues about changes. Lacking a single source of truth, engineers and mechanics spent valuable time checking and rechecking plans before project completion.

Another challenge was finding an effective strategy for rapidly rolling out and managing IT for users around the world. Running a global business, a big part of AWL's success would hinge on employees' ability to share large files, collaborate in real time, and communicate via video conferencing. AWL's IT team needed to stay lean and cost-effective even as it kept up with the rapid growth of users' IT maintenance and software requirements.

“Without NVIDIA virtual GPU technology, we would need to double in size and have local IT people on every site. We now support 700+ people with an IT team of 10 in the Netherlands and no local IT.”

Rody Kossen
Senior Systems
Engineer
AWL-Techniek

SOLUTION STATEMENT:

A solution came when AWL merged with the IT of MechDes in 2014, effectively inheriting the company’s IT expertise. As a firm that specialized in 3D SOLIDWORKS design, MechDes had already solved the problem of file sharing and high-performance graphics with Citrix and NVIDIA. Rody Kossen, who started at MechDes as a mechanical engineer, described the selection process, “When we looked toward virtualization in 2013, we discovered that Citrix was the only solution that delivered NVIDIA virtual GPUs. No one else provided that kind of high-level 3D graphics performance. Because NVIDIA was so advanced and provided so much flexibility, we had to go with Citrix.”

After the merge, Kossen transitioned to the role of senior systems engineer at AWL. As the company set out to create a roadmap for a global IT strategy, it leveraged Kossen’s extensive virtualization experience. During the next year, AWL built a global virtualized environment that would deliver high-performance graphics applications, purchasing Dell servers and installing NVIDIA GRID K2 and NVIDIA Tesla M60 GPU cards.

First, the team deployed a production environment on 12 servers with 150 concurrent NVIDIA virtual GPU (vGPU) users. When testing was complete, the IT team decided to allocate resources according to three profiles: engineer, viewer/mechanic, and office worker. To meet the resource requirements of engineers and mechanics, the IT team compared the performance of the NVIDIA GRID K2 and Tesla M60 cards. They discovered that they could get faster results and double the density per server by using the M60 GPUs (eight engineers on GRID K2 versus 16 on Tesla M60 and 16 mechanics on K2 versus 32 on M60).

AWL also discovered the value of the NVENC encoder built into NVIDIA GPUs in terms of delivering better user performance and data transfer. “We leveraged NVENC to reduce CPU usage on our server farm. We found that by using NVENC we could significantly reduce the performance impact on a host machine’s GPU and CPU. We could use one core less per user because the workload is offloaded to the GPU. This also lowers the latency for the user because the NVENC could encode much faster than the CPU,” said Kossen.

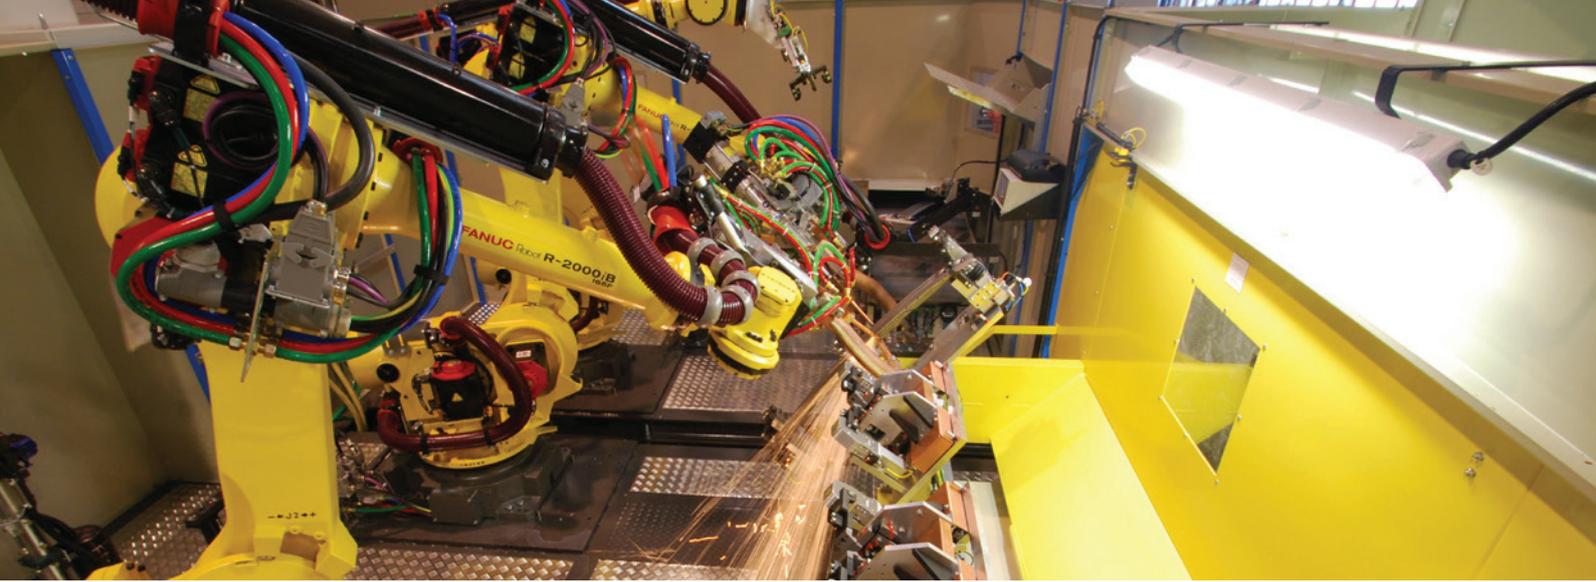


Image courtesy of AWL-Techniek

RESULT STATEMENT:

As AWL rapidly expands around the world, VDI powered by NVIDIA Quadro vDWS ensures the IT team can provide new virtual desktops to locations almost instantly. “Setting up a new location only takes a few hours. When we opened our U.S. location, we had staff working inside the office before the building was even ready. It was easy for our IT team to get everyone up and running remotely,” said Kossen.

As a result of deploying NVIDIA Quadro vDWS, the company has managed to almost double its engineering and mechanics staff without the need for additional IT support. “We have almost no IT equipment at our locations, so there’s almost no maintenance. Without NVIDIA virtual GPU technology, we would need to double in size and have local IT people on every site. We now support 700+ people with an IT team of 10 in the Netherlands and no local IT.”

NVIDIA Quadro vDWS enables engineers to work with more efficiency by providing unmatched 3D rendering speed and capacity. Moreover, engineers can work collaboratively on 3D graphics files in ANSYS, Siemens PLM NX, Dassault Systèmes CATIA, SOLIDWORKS, and Siemens Process Simulate. Mechanics can view updated files in eDrawings at any of AWL’s facilities without data synchronization issues. And user experience significantly improved—rotating models in SOLIDWORKS and opening the PLM Viewer and browsing projects became much faster.

With CPU offloading, engineers have a high-quality, fluid experience when video conferencing. This was particularly important as it’s not uncommon for engineers in the Netherlands to use Skype for Business and do screen sharing with engineers in China. “People are very happy with the solution because they can work in a fast and stable environment and are always working on the latest revision of files,” said Kossen.

“Citrix XenMotion support on NVIDIA vGPU is very critical, as we are a global company running almost 24 hours a day with a really short service window. Now I have the flexibility to take a host out of production without any issue. With today’s security issues—like Meltdown and Spectre—this is a ‘must-have’ to keep our environment patched and safe!”

Rody Kossen
Senior Systems
Engineer,
AWL-Techniek

High-performance virtualization also dramatically speeds up the production process. “Today our shop floor is 90 percent digital. You won’t see any paper plans or manuals. Mechanics use 3D viewers to view engineering plans and then build equipment directly from them,” said Kossen.

This improvement in productivity also extends to office staff using Microsoft Office as well as Skype for Business VoIP calls, video conferencing, and instant messaging. “Users can now work anywhere, anytime, with consistent performance. They work faster because they can use any device that they like and yet still enjoy great performance. That’s because of NVIDIA virtual GPU technology.”

Other benefits AWL has realized due to NVIDIA Quadro vDWS have been improvements in user density on servers, as well as decreased latency. “We saw user density increase 300 percent as we were evaluating the latest NVIDIA Pascal™ architecture-based GPUs (eight engineers with [NVIDIA] Kepler™ to 24 with Pascal), with at least the same performance per user. The total end-to-end latency dropped by 20 milliseconds when leveraging NVENC because the encoder is much faster. This resulted in a 10 percent decrease of the total end-to-end latency. On our Dell servers with NVIDIA Tesla M60 GPUs, we can host 16 people; with NVENC, we saw a total of 28 percent decrease of CPU usage (one core per user).” Now with the Dell R740 with NVIDIA Tesla P40 GPUs, they can even achieve 24 engineers or 48 mechanics on one server. Thus, by upgrading from K2 they can triple their capacity in the same physical space.

As AWL continues to expand, innovations from each iteration of NVIDIA vGPU software—like support for live migration—become invaluable. In the past, IT needed to log off all users from their Citrix environment when they wanted to perform maintenance on the hosts. In the beginning, when they had just a few users, this wasn’t a big issue. But as the company began to grow, it became a challenge. “Citrix XenMotion support on NVIDIA vGPU is very critical, as we are a global company running almost 24 hours a day with a really short service window. Now I have the flexibility to take a host out of production without any issue. With today’s security issues—like Meltdown and Spectre—this is a ‘must have’ to keep our environment patched and safe!” said Kossen.

Upcoming projects for AWL's IT team include opening a new headquarters in the Netherlands and migrating all users to Windows 10. Said Kossen, "We're creatively exploring new ways to make it easier for users to collaborate in our virtualized environment. Part of that is getting everyone on Windows 10, and another part is giving them the option of working side-by-side at floating workstations. We envision an office where engineers and mechanics can meet, log in to their profiles, and share ideas freely. NVIDIA virtual GPU technology ensures that, no matter what application they're working on or where they're located, staff can experience great performance."

Going forward, as AWL expands and acquires more companies, they'll continue to leverage NVIDIA virtual GPU technology on Pascal architecture-based cards to provide their employees with the best performance and IT with simplified management and cost-effective scalability.

To learn more about NVIDIA Quadro vDWS, visit: www.nvidia.com/quadro-vdws

www.nvidia.com/virtualgpu



© 2018 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, NVIDIA GRID, Kepler, Pascal, NVIDIA Quadro, and NVIDIA Tesla are trademarks and/or registered trademarks of NVIDIA Corporation. All company and product names are trademarks or registered trademarks of the respective owners with which they are associated. JUN18

