



Sponsored by: **NVIDIA**

Authors:

Linn Huang
Shannon Kalvar
Harsh Singh

May 2019

Business Value Highlights

434%

five-year ROI

7 months

to payback

49%

reduced total cost of VCC operations

36%

improved application performance

13%

more productive virtual desktop users

23%

faster full-time employee onboarding

51%

more efficient help desk staff

36%

reduction in costs for providing desktops

94%

reduction in unplanned downtime

NVIDIA Is Helping Organizations Provide Optimized Virtual Client Computing for Graphics and End-User Computing

EXECUTIVE SUMMARY

Virtual client computing (VCC), including application and desktop virtualization, is a foundational capability for organizations adapting to the future of work. The application of computing accelerators — including GPU acceleration — improves client performance, expands the available use cases, and provides greater resilience in technical performance. These in turn extend and preserve the business value of the VCC infrastructure investment.

IDC interviewed organizations that are using NVIDIA virtualization solutions (NVIDIA), which includes the NVIDIA Quadro Virtual Data Center Workstation (Quadro vDWS) for advanced, graphics-intensive workloads and NVIDIA GRID Virtual PC (GRID vPC) for mainstream virtual desktop infrastructure (VDI) use cases. On the basis of the data gathered, IDC found that these organizations are realizing significant benefits by leveraging NVIDIA's capabilities to deliver a higher-quality user experience for both enterprise and graphics-heavy applications. According to IDC's calculations, the organizations surveyed will realize annual benefits worth \$2.16 million per organization (\$446,000 per 100 devices) by:

- Driving improved performance and graphics capability for a variety of applications delivered to users via virtual desktops
- Improving the overall productivity of end users who rely on virtual client computing capabilities by enabling high performance for graphics-intensive workloads and an enhanced experience across all applications and workloads
- Optimizing capex and opex costs involved in providing access to applications to globally distributed users
- Reducing the incidence of unplanned downtime and minimizing help desk support requirements

SITUATION OVERVIEW

Virtual client computing, including application virtualization and desktop virtualization, is a foundational technical capability in modern operating environments. Most organizations have deployed some degree of these technologies to support a variety of use cases over the years, generally around the enduring themes of compatibility, mobility, and/or security requirements. During deployment and operations of this capability, companies invest in a combination of hyperconverged infrastructure systems, accelerators, and low-latency networks.

This foundational capability will have to transform as the nature of work evolves, such that by 2021, IDC predicts that at least:

- 60% of G2000 companies will actively monitor and manage employee experience (EX) and utilize EX as a key differentiator to build and maintain B2B and B2C relationships.
- 60% of G2000 companies will have adopted a future workspace model — a flexible, intelligent, and collaborative virtual/physical work environment — to improve employee experience and productivity.
- 65% of G2000 companies will respond the global demand for top talent by offering coworking and remote work options that leverage pervasive access to enterprise and collaboration applications.

IDC predicts that by 2024, 50% of structured repeatable tasks will be automated, and 20% of workers in knowledge-intensive tasks will have artificial intelligence (AI)-infused software or other digitally connected technology as a “coworker.”

To enable this, transformation organizations will have to adopt a continuous stream of improvements of the virtualization solution. This need for continuous capability escalation and performance is driven by multiple factors including:

- The ever-increasing capabilities, 3D, AI, and so forth in applications
- Increasing graphics requirements for operating systems and applications
- Increased processing power in endpoint devices and thin clients
- Increasing use of multiple monitors and high-resolution monitors
- Interactive 2D and 3D websites, using APIs such as WebGL
- The density of data transmission/utilization

- Larger model sizes and data sets that employees work with and need to collaborate on
- Increase in video content and video collaboration

This dynamic will only accelerate as capabilities including deep learning/inference on the edge, IoT coordination, and application disaggregation are included into application and operating system software.

The state of IT budgeting further complicates this dynamic. Projects are planned for multiple phases, each delivering a specific value, and then only the first phase is funded. Years later the installed systems have degraded over time and need to be replaced, usually by another multiphased project where only one phase will receive sufficient funds. The shift to zero-balance budgeting will likely extend this trend leading to additional short-term thinking and quarter-by-quarter manipulations on systems and changes that require continuous work over an extended period.

NVIDIA GPU-ACCELERATED VIRTUALIZATION SOLUTIONS

The trends outlined previously — in budgets, technology, and the future of work — mark fundamental shifts in the concept of work as we know it. For enterprises that rely on virtual client computing, a new method of delivering virtual desktops is required, one that can keep pace with the digital workplace and advance with new technology developments. NVIDIA's virtualization solutions support virtual desktops and workstations environments with advanced AI and graphics capabilities, modernizing virtual desktop infrastructure for the digital age. NVIDIA virtual GPU (vGPU) software enables delivery of graphics-rich virtual desktops and workstations accelerated by NVIDIA GPUs, on the market today. With NVIDIA virtual GPU software, GPU resources are shared across multiple virtual machines (VMs), or multiple NVIDIA GPUs can be allocated to a single virtual machine to power the most demanding workflows. NVIDIA virtual GPU software runs on NVIDIA GPUs such as the T4 and M10 as well as the Quadro RTX 6000 and Quadro RTX 8000.

NVIDIA offers three software products suited for enterprise organizations. NVIDIA Quadro Virtual Data Center Workstation is designed to meet the performance requirements for designers, architects, engineers, and artists. When paired with a powerful NVIDIA GPU, users can virtualize any application from the datacenter with an enhanced user experience — including ANSYS Discovery Live, Esri ArcGIS Pro, Siemens NX, Dassault Systèmes SOLIDWORKS, and Autodesk Revit — allowing workstation-class performance on any device.

NVIDIA GRID Virtual PC is designed for the modern digital worker in virtual desktop environments. The term digital worker refers to the way we work in the tech-savvy, digital age. Across industries, the workplace has become more visual — video, web conferencing, multimonitors, dynamic browsers, 3D features in everyday applications, and more are the new normal. Even Windows 10 is bringing more and more 3D capabilities to base operating systems, further reinforcing the need for GPUs with NVIDIA GRID vPC software to deliver an accelerated virtual graphics experience.

NVIDIA GRID Virtual Apps (GRID vApps) is used to launch applications on any device without having to present a full, virtualized desktop to a user. Remote desktop session host (RDSH) solutions can be paired with a more powerful GPU to run more graphics-intensive applications or paired with a less powerful GPU to run general-purpose applications and have more users share a virtual machine.

THE BUSINESS VALUE OF NVIDIA SOLUTIONS

Study Demographics

IDC conducted research that explored the value and benefits for organizations of using NVIDIA virtual GPU solutions, including Quadro vDWS, GRID vPC, and GRID vApps, to provide access to virtual client computing across their organizations. The project included 10 interviews with organizations using NVIDIA vGPU solutions that have detailed experience with the benefits and costs of using them. During the interviews, these companies were asked a variety of quantitative and qualitative questions about the impact of deploying these NVIDIA solutions and their ability to deliver applications and workloads via VCC, as well as VCC costs and staffing requirements.

Table 1 presents the study demographics for the companies surveyed. The organizations interviewed have an average employee size of about 42,000. Almost every one of these employees is utilizing IT services and are being supported by an IT staff of about 816 employees. These organizations were running about 528 business applications, 16 of which were high-end professional 3D applications. The vertical industries being represented include manufacturing, professional services, financial services, energy, government, healthcare, and education sectors. *(Note: All numbers represent averages.)*

TABLE 1 Study Demographics

Factor	Average	Median
Number of employees	42,843	4,900
Number of IT staff	816	36
Number of IT users	41,796	1,500
Number of business applications	528	100
Number of business applications requiring graphics	16	14
Organizational revenue	\$51 billion	\$2 billion
Industries	Manufacturing (3), professional services (2), financial services, energy, government, healthcare, and education	

n=10 Source: IDC, 2019

Organizational and Selection Use of NVIDIA Solutions

IDC interviewed 10 NVIDIA vGPU customers, with 6 of those customers primarily using Quadro vDWS, 3 customers mostly using GRID vPC, and 1 customer using both products extensively. Study participants described company usage patterns for the respective NVIDIA solutions as well as providing a snapshot of overall IT and business environments supported. Most organizations interviewed moved to NVIDIA from traditional desktop environments, with remaining migrations to NVIDIA from legacy VDI solutions that were based on CPU only. Interviewed NVIDIA customers discussed both the value proposition and the rationale behind their choice of NVIDIA solutions, with the drivers differing to some extent by NVIDIA product.

Overall, study participants discussed several common areas of perceived value to both NVIDIA products in delivering VCC capabilities, including:

- The ability to handle large data sets
- Managing business and operational growth without increasing overall footprint
- Cost savings as a result of virtualizing resources
- Being able to clone “on the fly” and use linked pools
- Single pane of glass desktop management
- Improved collaboration with increased security and version control
- Improved user experience and employee engagement

However, they also discussed distinct evaluation criteria. For respondents using the NVIDIA Quadro vDWS, the primary concern was to provide a VCC solution sufficiently robust to run and ensure high performance for very graphics-intensive applications and workloads. Study participants uniformly discussed having specific performance-related requirements for these workloads that any VCC solutions had to meet. They chose NVIDIA after concluding that NVIDIA Quadro vDWS was the VCC solution best suited to deliver high-quality experiences for these types of graphics-heavy applications and workloads:

- **Supporting graphics-intensive applications globally:** *“We switched to NVIDIA because we have a set of applications that are fairly graphics intensive. We have users around the world who are using these applications running on a centralized datacenter. So we needed a way to have everybody be able to connect remotely while having good graphics capabilities.”*
- **Managing growth without increasing footprint:** *“We have been growing substantially by doubling our student population over a period of 10 or 11 years. Our physical square footage was becoming too small, so we saw virtualization as a solution. The cost of expansion was prohibitive and required us to build new buildings. Part of the solution involved consolidating our operations using Quadro vDWS to supplement our existing virtualization solution. Diversifying our computer labs by virtualizing some of them freed up space for classrooms. This saved \$250,000 per classroom that we would have had to build otherwise.”*

Organizations using GRID vPC spoke of the performance advantages such as allowing users to access key data sets more easily but also considered broader efficiencies among IT staff and end users. They discussed factors that they considered in choosing GRID vPC to support knowledge workers across their organizations:

- **More efficient desktop duplication:** *“We needed a better way to manage our environment with a lean IT staff. Having everybody crawling under desks and building machines the old way was time consuming. We also wanted the ability to clone on the fly, the ability to use linked pools, and a single pane of glass for viewing all desktops.”*
- **Better VDI experiences:** *“We looked at two different vendors for our VDI platform, and they were not at the same place that NVIDIA was. There was no competition to what NVIDIA could do and our users would not be happy with VDI without NVIDIA. It was too slow and graphics rendering was poor at best. NVIDIA gave us a way to use VDI the way we needed to and gain full advantage of our virtual infrastructure.”*

Table 2 shows that the average number of users of both GRID vPC and Quadro vDWS solutions was substantial at 209 and 626 users, respectively. 76% of those users were full-time employees and 6% were part-time users, with contract workers making up most of the remaining 18% of users. Additional usage patterns are presented in Table 2.

TABLE 2 Organizational Usage of NVIDIA vGPU

NVIDIA Product Use	Average	Range
GRID vPC user base	209	12–1,200
Quadro vDWS user base	626	45–5,100

Source: IDC, 2019

Table 3 provides granular data on the IT environments that characterized surveyed companies as well as specifics on how NVIDIA solutions were being deployed. For example, the average number of devices was 486, while the average number of applications was 35. In terms of device profiles, 32% were desktop devices, while 41% were laptops and 26% thin clients.

TABLE 3 Organizational Environment of NVIDIA vGPU

NVIDIA Product Use	Average	Range
Number of offices	9	4
Number of devices	486	313
Number of applications	35	20

Source: IDC, 2019

Business Value Analysis

IDC’s analysis found that interviewed organizations were using NVIDIA solutions to provide the levels of virtual desktop performance and enhanced graphics capability that they needed to support their growing business operations. IDC found that deployment of the NVIDIA solutions offered significant performance and operational benefits related to delivering virtual client computing including:

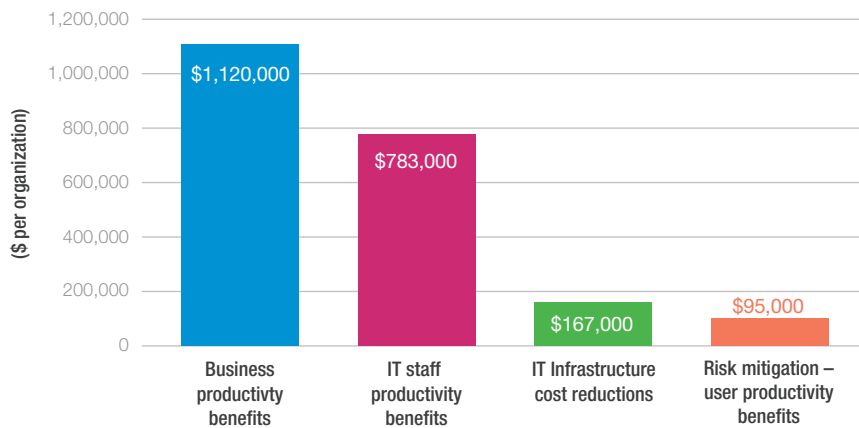
- 36% improved application performance
- 13% more productive virtual desktop users
- 51% more efficient help desk operations

These benefits generated positive ripple effects throughout the organizations surveyed, resulting in reduced IT capex and opex spending and decreased unplanned downtime. Using the survey data collected from all 10 companies, IDC quantified these benefits further:

- 36% reduction in costs for providing desktops
- 434% five-year ROI
- 49% reduced TCO

On the basis of IDC’s calculations, the NVIDIA customers surveyed realized average annual benefits worth \$2.16 million on a per organization basis (\$446,000 per 100 devices) as shown in Figure 1.

FIGURE 1 Average Annual Benefits per Organization



Source: IDC, 2019

Study participants provided detailed comments about these benefits:

- **The ability to manage desktops remotely (GRID vPC customer):** *“Maintaining clients in a dispersed environment is a manpower issue. It also introduces a lot of variabilities because you might have inconsistency among computers. A VDI environment totally obviates that need to spend more time on support, which is why we went to VDI and we get the same benefits from NVIDIA vGPU for graphics workstation support.”*
- **Better flexibility and management (Quadro vDWS customer):** *“There is more flexibility such as changing the configuration of a virtual desktop as needed or spinning up a different configuration as needed ... It makes it that much easier for us to manage as a small company that does not have the typical management tools that a larger company has. It allows us to manage one or two images and have consistency across the entire organization.”*
- **Consistent desktop environments across the organization (Quadro vDWS customer):** *“Compliance is one of the most significant benefits ... making sure the experience is consistent*

across all virtual workstations. Virtualization has allowed us to maintain multiple customer design environments in a completely virtual way for different customers. We would spend lots of time on managing those physical environments that we no longer have to do anymore with VDI."

While IDC did not find any significant difference in the value that interviewees with Quadro vDWS or GRID vPC users were achieving, those organizations that utilized these respective products did describe to IDC the key benefits these specific products were adding to their organizations, which is discussed in the next two sections.

Differentiated Value of Quadro vDWS

Study participants spoke with IDC about the critical importance of NVIDIA Quadro vDWS being able to provide high quality and high performance via VCC for graphics-heavy applications and workloads. They noted needing to be able to provide access to these applications to support and enhance distributed operations of their companies but struggled to accomplish this before deploying Quadro vDWS.

In particular, technical professionals in these companies who work with graphics-intensive applications depend on immersive visual computing for a variety of applications including interactive design processes and performing real-time simulations. Customers reported that Quadro vDWS has allowed these users to move beyond the limitations of high-powered workstations tethered to physical facilities by accelerating and enhancing graphics applications and still providing all the other benefits of VCC technology.

These organizations added additional context for the benefits they were seeing with Quadro vDWS:

- **Ensure the performance of data-heavy graphics applications:** *"We chose Quadro vDWS because we are running engineering applications that require discrete graphics. We tried another approach to VDI, and the performance was barely usable. NVIDIA vGPU has turned out to be the best fit for our organization because it meets both our density and our performance requirements."*
- **Enable the use of VDI for graphics-heavy workloads for global workforce:** *"We have CAD designers who are not in our main headquarters who complained about slow performance in opening models from our key databases. We deployed Quadro vDWS instead of having expensive workstations. Without Quadro vDWS, we wouldn't be able to do this with traditional CPU-only VDI."*

- **Provide performance to support remote contractors:** *“The overall business driver for choosing Quadro vDWS was being able to have contract workers do CAD work offsite — NVIDIA allowed that to happen because it was previously much more difficult to support CAD users remotely, and Quadro vDWS makes that possible. It’s a regulatory issue for us: If we want to use contractors, we cannot have them in our facilities, and we lower expenses.”*

Differentiated Value of GRID vPC

The benefits that NVIDIA GRID vPC provided included the ability to expand the number of applications that could be enabled on their VDI environment with access to the application’s full functionality (e.g., advanced features only supported with GPUs in Adobe Photoshop and Google Earth). Some participants stressed how GRID vPC enabled better user experience for everyday applications such as Microsoft Office, Google Chrome, and PDF viewers. Interviewees also cited good support for the use of corporate-level video applications for training and other purposes. Study participants expanded on these benefits:

- **Could not run any graphics application without NVIDIA:** *“When we decided to do VDI, we had to be able to offload things like video. NVIDIA GRID vPC enabled us to do that and applications like AutoCAD (2D) and Adobe Photoshop. We couldn’t do VDI without GRID vPC because you can’t have any graphics processing since it’s all done in CPU cycles on servers, which does not run well.”*
- **Increased application use on VDI-enabled cost savings:** *“There were two different levers desired. In the case of the first lever, we had a very good functional VDI that did not include a GPU, but we wanted to expand the number of applications that could be supported, such as Photoshop and Google Earth with plug-ins. As application performance improved, the second lever eventually became an opportunity to replace high-cost workstations being deployed for end users.”*
- **NVIDIA — the complete solution:** *“We performed an analysis to determine if there is an alternative. In the case of vGPU, NVIDIA really is the only vendor that has a complete solution stack. We looked at our options for software and some light hardware augmentations as opposed to a fairly complete hardware solution to graphics processing, but at the end of the day, there really was no alternative to NVIDIA.”*

Business and End-User Benefits and Impacts

NVIDIA customers benefited from the fact that vGPU solutions, including NVIDIA Quadro vDWS and GRID vPC, provided the levels of performance that their organizations required for growing and supporting their business operations. In a virtuous cycle, better performance and reliability

led to the end result of developing and supporting better products and services, better customer service, and the ability to move offerings to market more quickly.

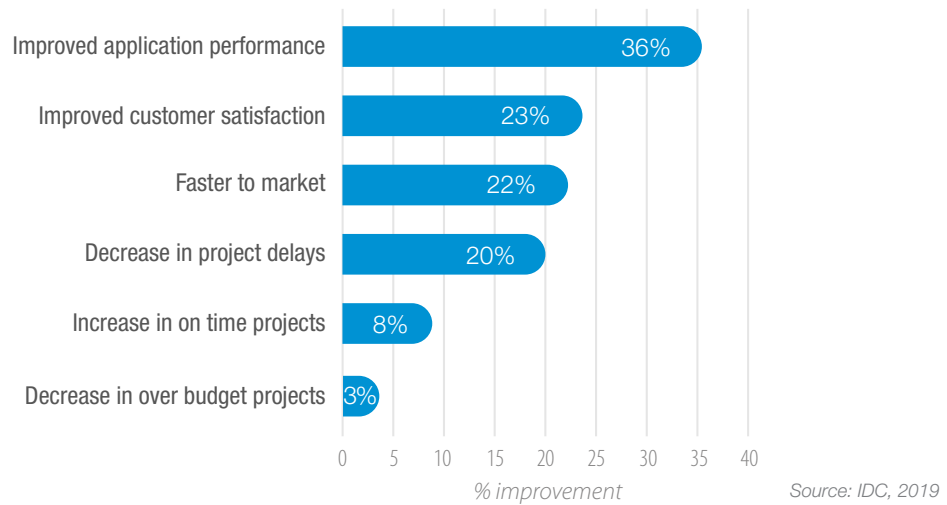
The end result for these NVIDIA customers has been strong virtual client computing performance that delivers operational efficiencies in the form of higher levels of employee productivity and increased revenue through business enablement. As described previously, organizations are using Quadro vDWS and GRID vPC for different workloads, but the core benefits that led to these improvements are similar. Study participants cited providing employees with the ability to work remotely while being able to access any application — from 3D CAD and other high-performing visualization applications to business productivity and multimedia-rich applications, increased mobility, and greater availability of corporate resources for remote users. The process of employee onboarding also recognized improvements. Study participants discussed these benefits in detail:

- **Greater confidence for end users (Quadro vDWS):** *“Employees can know the environment they are working on is up to date. Any new features or bug fixes have been integrated into that virtual environment so that compliance works.”*
- **End users to access data and apps anywhere (Quadro vDWS):** *“An important piece of doing VDI with NVIDIA is the access anywhere. Employees can work at home, work offsite, or work when they travel. Prior to VDI they couldn’t do that; they didn’t have access to the CAD environment.”*
- **Better flexibility for users (Quadro vDWS):** *“The agility we get with NVIDIA vGPU is understood by the business. The biggest thing it gives our users is the ability to work anywhere. People can work at home and don’t have to stay in the office to meet deadlines and get work done. You can get your work done wherever and whenever you want.”*
- **Work from anywhere (GRID vPC):** *“We wanted to have somewhat open offices, in the sense that no one really has their own office. Everyone sits in the same area where they can log in with the same monitors and same zero clients. You can walk into the same position at any floor or in any country and then log in as a known user. We have gained the flexibility of not having users tied to desks.”*
- **Greater mobile usage supported (Quadro vDWS):** *“The biggest benefit is for our students. They can work on their own timetable in their own space and don’t have to abide by rules that might say you need to be in this room ... It also enables students who might have shied away from STEM tracks because of the extra costs needed like a big \$3,000 laptop since they can now access the same application on a \$250 Chromebook.”*

Figure 2 shows the business impacts experienced by NVIDIA customers. Application

performance improved substantially (36%) as did customer satisfaction (23%). Other quantified benefits included moving products and services to market faster (22%) and decreasing the number of project delays (20%).

FIGURE 2 Business Impact Metrics



IDC also measured these benefits as they applied to end users. As shown in Table 4, an average total of 835 end users gained 36 productive hours annually on a per-user basis as the result of efficiencies gained through access to virtual client computing for specific applications and improved performance. When this benefit is translated into financial terms (the value of end-user time), the overall average gain is substantial at \$1.12 million per year. The impact in terms of FTE equivalence annually was 16.

TABLE 4 End-User Impact

Enhanced User Productivity	Per Organization
Number of users impacted	835
Productive hours gained per year	30,000
Productive hours gained per user per year	36
End-user impact (FTE equivalent per organization per year)	16
Value of end-user time per year	\$1.12 million

Source: IDC, 2019

Another business benefit is related to improvements in employee onboarding related to greater ease in providing new employees with high-quality access to applications and workloads. IDC quantified these benefits for both full-time regular employees and contract employees. Before the NVIDIA vGPU deployment, full-time regular employees were onboarded in 9.8 days. After deployment, the process took 7.5 days, a 23% level of improvement. Even higher levels of improvement were recognized for contract employees (33%) (see Table 5).

TABLE 5 Employee Onboarding

	Before NVIDIA	With NVIDIA	Difference	Change (%)
Time to onboard new full-time employees (days)	9.8	7.5	2.3	23
Time to onboard new contract employees (days)	0.8	0.5	0.2	33

Source: IDC, 2019

Increased Efficiencies for IT Operations

Study participants reported that NVIDIA vGPU solutions, including NVIDIA Quadro vDWS and GRID vPC, supported and improved significant aspects of their IT environments and enhanced the work of operational teams. These benefits centered on a number of factors:

- Ease of managing security and other required updates
- Reducing the overall cost of providing desktops
- Eliminating over- or under-allocation of resources with the ability to rightsize VMs
- Reducing the number of issues requiring remediation
- Ability to deal with problems — proactively, more quickly, and effectively

Traditional approaches to managing PC resources generally require more care and feeding in terms of staffing and time-consuming touch labor. In addition, there is a lot of variability and inconsistency among computers in terms of routine tasks that need to be performed such as patching. Deploying VDI environments avoid these issues for IT support teams including help desk. Study participants provided specifics on these benefits:

- **Ease of IT management (GRID vPC):** *“From an IT perspective, the NVIDIA vGPU solution makes control of the baseline and regimentation of how updates are done with security in a very consistent manner across the environment. That becomes inherited from your large VDI protocols. You don’t have these one-off desktops managed as a stovepipe because you*

now have centralized infrastructure managed as a unit. So from a configuration consistency standpoint, including things like patch quality and overall security posture — all those checklist items — it gives you more control and regimentation.”

- **Easier to fix issues (Quadro vDWS):** “Our unique issues have decreased, and it is much easier to resolve them now. If somebody has an application issue, we can spin up an identical desktop without bothering them. We don’t have to sit at their machine or take their machine. We can also work on the problem on a separate machine. We can detect problems, and then prevent them from occurring for other people as well. The majority of the issues we get are not unique to one machine, they are often issues that others will have eventually. So there is a preventive element to this as well.”

Table 6 presents the quantified benefits for end-user device management impacts. The management of end-user devices saw a 25% level of improvement. This and other improvements resulted in a substantial 25% cost savings on an annual basis.

TABLE 6 End-User Device Management Impact

	Before NVIDIA	With NVIDIA	Difference	Change (%)
Management of end-user devices	2.4	1.8	0.6	25
Hours per device per year	9.3	6.9	2.3	25
Staff time cost per year (\$)	239,048	179,107	59,941	25

Source: IDC, 2019

Impact on Unplanned Downtime

NVIDIA customers reported positive impacts with respect to unplanned downtime and overall business productivity. Surveyed companies were able to reduce the incidence of outages after deploying Quadro vDWS and GRID vPC solutions. This benefit also positively affected the line-of-business (LOB) users and business outcomes.

Table 7 shows the unplanned downtime impacts for the NVIDIA platforms. The frequency per year showed a substantial level of improvement (75%). In addition, the mean time to repair (MTTR) showed a 63% level of improvement. Especially noteworthy is the fact that FTE productivity impacts showed a 94% level of improvement as did the annual cost of unplanned downtime.

TABLE 7 Unplanned Downtime Impacts

	Before NVIDIA	With NVIDIA	Difference	Change (%)
Frequency of incidents per year	16.9	4.2	12.7	75
MTRR (hours)	4.8	1.8	3.0	63
Productivity impact (FTE equivalent)	1.6	0.1	1.5	94
Hours per device per year	6.1	0.4	5.7	94
Cost of unplanned downtime per year (\$)	110,445	7,074	103,372	94

Source: IDC, 2019

Table 8 shows unplanned downtime impacts extended to help desk operations within the organizations surveyed. For example, the number of calls per week was significantly diminished (46%). Annualized help desk FTEs showed a 51% level of improvement and MTRR improved by 28%. These benefits translate to financial impacts and staff time savings of \$331,504, a 51% level of improvement.

TABLE 8 Help Desk Impacts

	Before NVIDIA	With NVIDIA	Difference	Change (%)
Calls per week	281	151	130	46
MTRR	6.3	4.5	1.8	28
Help desk FTEs	6.6	3.2	3.3	51
Hours per device per year	25.4	12.5	12.8	50
Staff time cost per year (\$)	655,207	323,703	331,504	51

Source: IDC, 2019

More Cost-Effective Operations

Study participants spoke with IDC about how NVIDIA vGPU platforms were designed to be a cost-effective solution for providing desktop resources and the variety of applications they support. They discussed how using a VDI solution allowed them to pare down both capex and opex associated with more traditional PC resources.

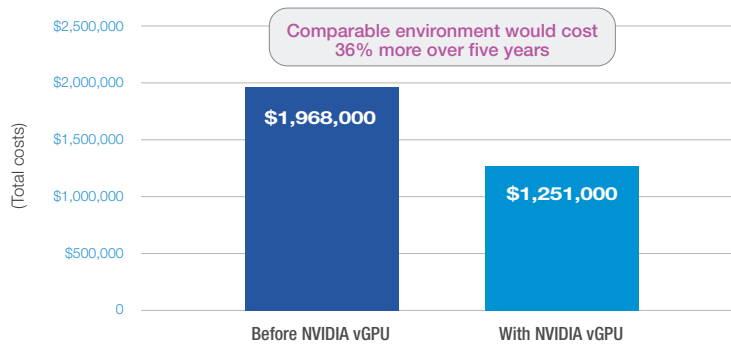
This was the case for a number of reasons including the fact that companies were able to avoid buying new hardware as their business needs expanded or evolved. One study participant spoke about NVIDIA vGPU in terms of these benefits: “We projected that it would be easier and less expensive to go with NVIDIA vGPU instead of refreshing workstations in a

staggered manner and having people on different generations of desktops. Our projection was that we could save up to 30% overall.”

Figure 3 presents the desktop savings associated with the deployment of NVIDIA solutions. Given that 85% of companies surveyed had relied on legacy desktop environments prior to deploying NVIDIA, IDC calculated that NVIDIA customers were able to realize a 36% savings over alternatives over a five-year period.

and performance that they can translate through to improvements to core business results and how their organizations operate. On average, these organizations attributed increased revenue worth \$3.54 million per year to their use of Red Hat solutions (\$12,600 per 100 users) by helping them better address and win business opportunities.

FIGURE 3 Five-Year Desktop Costs




Source: IDC, 2019

These cost savings also extended to opex. IDC calculated the cost of operations over a five-year period. As shown in Figure 4, costs shifted from \$8.21 million to \$4.20 million over the period, thereby lowering overall costs substantially (49%). IDC used three evaluation criteria to arrive at these calculations:

- Cost of lost productivity and unplanned downtime
- Costs associated with IT staff time including management
- Costs of NVIDIA vGPU deployment versus alternative approaches

FIGURE 4 Five-Year Cost of Operations



Source: IDC, 2019

ROI Analysis

Table 9 presents IDC’s analysis of the benefits and costs related to surveyed organizations’ use of NVIDIA solutions. IDC projects that these organizations will, over five years, realize discounted benefits of \$7.6 million per organization (\$1.5 million per 100 internal users). When compared against a discounted investment of \$1.4 million per organization (\$0.3 million per 100 internal users), these organizations will see an ROI of 434% and a breakeven on their investment in seven months.

TABLE 9 Five-Year ROI Analysis

	Total	Per 100 Devices
Benefit (\$M)	7.6	1.5
Investment (\$M)	1.4	0.3
Net present value (NPV)	6.1	1.2
ROI (NPV/investment) (%)	434	434
Payback (months)	7	7
Discount factor (%)	12	12

Source: IDC, 2019

CHALLENGES AND OPPORTUNITIES

Both application and client virtualization software have historically addressed narrow use cases around compatibility, mobility, and security supported by a combination of specialized infrastructure and low-latency networking. As discussed in the Situation Overview section, this system has proven to be expensive to operate over time because of a combination of technical, innovation, and budgeting factors.

The stability of these core use cases, though, presents a unique opportunity for organizations that can apply lateral thinking to their infrastructure investments. Traditionally, there is always another use for low-latency networking and the hyperconverged infrastructure. GPU acceleration in the datacenter, though, has traditionally been more difficult to apply to alternative use cases.

That has changed in recent years. GPU accelerators, based on parallel processing architecture, are well suited for a number of artificial intelligence training and inference models. They can also provide, through desktop and application virtualization, the ability to support graphically intensive applications on lower-capability devices, expanding business platform use cases in the traditional environment. They can also support innovation projects in the augmented reality and data analytics fields, which would otherwise go under resourced.

The relative simplicity of budgeting for and supporting GPU acceleration as part of a project, relative to hyperconverged systems and/or low-latency networks, also factors into the long-term return for virtual client computing. By providing a consistent performance boost through a single investment, these systems stabilize ROI over time. Hyperconverged systems require regular updates and constant operational maintenance/tuning, while low-latency networks require the same in addition to intensive partner management and data management capabilities.

There remains an open question as to whether accelerated virtual client infrastructure is better provided through a public cloud or a private datacenter. It is likely that public cloud provisioning will become increasingly viable over time, especially for the core use case of mobility and emerging use cases around the future of work. However, as IT begins to use infrastructure to support lateral innovation use cases, it may make more strategic sense to own and manage the underlying assets. This consideration will be especially important over the next three to five years as an increasing number of business-critical applications move from private datacenters to the cloud.

This dynamic will lead the deployment of “hybrid cloud” (multiple public and private cloud infrastructures) for the foreseeable future. Although not directly addressed in this study,

public cloud components of a hybrid cloud VCC infrastructure will likely also benefit from computing accelerators, including GPU acceleration. This trend is already visible in the client workload offerings from the major public cloud vendors, which include GPU-accelerated options.

SUMMARY AND CONCLUSION

The NVIDIA accelerated datacenter solution set provides an estimated 49% decrease in the operating costs of virtual desktop infrastructure over a five-year period, increases the number of use cases that infrastructure can support, and provides improved performance in modern, graphically intensive applications. These benefits are complicated by industrywide shifts from private datacenters to hybrid public/private computing infrastructures and continuing issues in IT budgetary models, which make the long-term value of large infrastructure investments difficult to sustain. This complication does not prevent the use of GPU acceleration for client virtualization solutions as they are available in both public and private forms.

APPENDIX

IDC's standard ROI methodology was utilized for this project. This methodology is based on gathering data from current users of NVIDIA vGPU as the foundation for the model. Based on interviews with organizations using NVIDIA vGPU, IDC performed a three-step process to calculate the ROI and payback period:

- **Gathered quantitative benefit information during the interviews using a before-and-after assessment of the impact of NVIDIA vGPU.** In this study, the benefits included staff time savings and productivity benefits and operational cost reductions.
- **Created a complete investment (five-year total cost analysis) profile based on the interviews.** Investments go beyond the initial and annual costs of using NVIDIA vGPU and can include additional costs related to migrations, planning, consulting, and staff or user training.
- **Calculated the ROI and payback period.** IDC conducted a depreciated cash flow analysis of the benefits and investments for the organizations' use of NVIDIA vGPU reports over a five-year period. ROI is the ratio of the net present value (NPV) and the discounted investment. The payback period is the point at which cumulative benefits equal the initial investment.

IDC bases the payback period and ROI calculations on a number of assumptions, which are summarized as follows:

- Time values are multiplied by burdened salary (salary + 28% for benefits and overhead) to quantify efficiency and manager productivity savings. For purposes of this analysis, based on the geographic locations of the interviewed organizations, IDC has used assumptions of an average fully loaded salary of \$100,000 per year for IT staff members and an average fully loaded salary of \$70,000 per year for non-IT staff members. IDC assumes that employees work 1,880 hours per year (47 weeks x 40 hours).
- The net present value of the five-year savings is calculated by subtracting the amount that would have been realized by investing the original sum in an instrument yielding a 12% return to allow for the missed opportunity cost. This accounts for both the assumed cost of money and the assumed rate of return.
- Because IT solutions require a deployment period, the full benefits of the solution are not available during deployment. To capture this reality, IDC prorates the benefits on a monthly basis, and then subtracts the deployment time from the first-year savings.

Note: All numbers in this document may not be exact due to rounding.

IDC Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-insights-community.com
www.idc.com

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2019 IDC.
Reproduction without written permission is completely forbidden.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.