



IBM SmartCloud Desktop Infrastructure for accelerated graphics

Solution guide using IBM iDataPlex and IBM NeXtScale servers

IBM Systems and Technology Group ISV Enablement Team

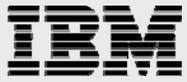
April 2014

 @IBMSystemsISVs



Table of contents

Introduction	1
Business value	1
Classes of users for graphics acceleration	2
Solution overview	3
Hardware platform	7
Graphics adapters.....	7
Compute servers.....	7
Networking	10
Shared storage	10
Solution architecture	11
Virtual Desktop Graphics Acceleration	11
Host Shared Graphics Acceleration.....	14
Citrix XenDesktop	15
VMware Horizon View.....	16
Reference design	17
Aircraft company example	18
Publishing company example	19
Ordering information	21
Resources	25
Trademarks and special notices	26



Introduction

The IBM® SmartCloud® Desktop Infrastructure (SDI) offers robust, cost-effective, and manageable virtual desktop solutions for a wide range of clients, user types, and industry segments. These solutions can help to increase business flexibility and staff productivity, reduce IT complexity, and simplify security and compliance. Based on a validated IBM reference architecture design, the infrastructure supports various hardware, software, and hypervisor platforms and ensures high level of confidence in component interoperability and workload support.

Increasingly, there is a customer demand for more sophisticated virtual desktops that support multimedia, business graphics, engineering graphics and technical computing visualization. This solution guide discusses the SmartCloud Desktop Infrastructure solutions for accelerated graphics for Citrix XenDesktop or VMware Horizon View using IBM iDataPlex® or ultra-dense IBM NeXtScale System servers with graphics processing units (GPUs). This solution guide only addresses Microsoft® Windows® users and not Linux® users.

Business value

Virtual desktop infrastructure (VDI) gives users the freedom to access their desktop applications from any device. It has recently developed traction in the market due to its many benefits, including increased security and reduced total cost of ownership (TCO). Companies that offer VDI to their employees reduce the number of desktop applications and licenses they must manage, and subsequently reduce their ongoing operational costs. At the same time, with fewer applications to manage, companies can more tightly control authorized access to the applications, protect corporate data, and ensure standards compliance.

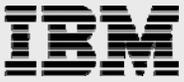
Several key factors drive virtual desktops in today's business climate.

- Data security and compliance concerns
- Complexity and costs of managing existing desktop environments
- An increasingly mobile workforce
- The changing ownership of end-point devices with bring your own device (BYOD) programs
- The need for rapid recovery from theft, failure, and disasters

Accelerated graphics using GPUs is needed to accelerate the rendering of 3D scenes created by computed-aided design (CAD), computer-aided engineering (CAE), petrochemical or other technical computing or financial applications and is also very useful for multimedia content creation, simple 3D business graphics, and games. Note that this is different from *accelerating computation* using general processing on GPUs (GPGPUs).

The factors that are driving the deployment of virtual desktops are even more important for users who need accelerated graphics. Here are some examples of the business value of VDI using GPUs:

Mechanical CAD diagrams consist of one or more large files that contain the information to draw an individual part, an assembly, or a complete structural design for a car, truck, or airplane. A user has to wait for all of these files to be downloaded to their CAD workstation before they can start any work and all of the updates have to be uploaded back into the file repository. With virtual desktops, the data only needs to be transported within the data center from the file repository to the virtual desktop server. This saves time and network bandwidth. For CAE and other kinds of technical computing, the data to be downloaded is often the result of a compute-intensive calculation and could be terabytes in size, which has to be stored locally before graphical images can be rendered on the workstation.



Centralization of data in the data center also helps to ensure its security. This means that intellectual property never leaves the data center (only rendered images are sent to clients). This allows vendors or partners to work on a shared design with a reduced risk of access to the complete design or theft of the intellectual property.

Centralization of work into virtual desktops in a data center can also afford an economy of scale where these desktops can be shared across workers in different time zones or with different requirements for access to accelerated graphics.

An employee on the manufacturing floor can easily bring up a view of the design on a wireless notepad or tablet to verify some aspect of the manufacturing process. Technical employees no longer need access to expensive workstations and use a simple thin-client machine with one or more screens to do their work perhaps while traveling or at a remote location. It also means that the same client machine can be used for general office work such as email, thereby eliminating a second client machine for the employee. In the case of a disaster, it is easy to setup client machines in a new location without the complexity of installing, configuring, and licensing a complete engineering workstation.

A financial analyst can have up to four high resolution displays of stock data and perform *what-if* computations without requiring the clutter and noise of a powerful workstation.

The SmartCloud Desktop Infrastructure approach offers the following advantages:

- Lowers the TCO over an extended period compared to traditional personal computers (PCs)
- Simplifies desktop administration, support, and management
- Enhances security and compliance management
- Improves availability and reliability
- Enables users to work anytime, anywhere quickly and easily, regardless of location or device
- Better supports growth initiatives for mobility and flexible work locations

Classes of users for graphics acceleration

A wide range of users might need access for accelerated graphics and their performance requirements vary from high-end engineering CAD users down to *knowledge workers* who occasionally need access to graphics acceleration.

For non-graphics acceleration, the terms of task worker, knowledge worker, and power user are often used. In this definition, task workers do not use graphics acceleration. Users who do require graphics acceleration can be split into three classes - GPU knowledge worker, GPU power user, and designer/engineer. Table 1 shows the types of applications that might be used for each class of user.

GPU knowledge worker	GPU power user	Designer/Engineer
<ul style="list-style-type: none">• Photoshop• 3D Viewer• Google Earth	<ul style="list-style-type: none">• Medium CAD• CAE visualization• Financial analyst	<ul style="list-style-type: none">• High-end CAD• Petrochemical seismic time reversal or reservoir analysis• Other technical computing visualization such as weather, life sciences

Table 1: Classes of users

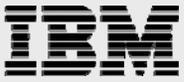


Table 2 compares the six classes of users and their typical requirements on virtual machine (VM) memory, virtual CPUs (vCPUs), display requirements, and GPU equivalent if they used a workstation. The Windows 7 Aero theme requires a vCPU and it is recommended to turn off this feature and instead use this valuable vCPU for extracting and compressing the rastered image from the GPU.

Class of user	VM memory	vCPUs	Resolution	Displays	GPU equivalent
Task worker	1 – 2 GB	< 1	1280 x 720	1	N/A
Knowledge worker	2 – 3 GB	2	1920 x 1080	1	N/A
GPU knowledge worker	2 – 4 GB	3	1920 x 1080	1	Quadro K600
Power user	3 – 6 GB	2	1920 x 1080	1	N/A
GPU power user	4 – 8 GB	3	2560 x 1600	1 or 2	Quadro K2000 or K4000
Designer/Engineer	8 – 16 GB	4 or more	2560 x 1600	Up to 4	Quadro K5000 or K6000

Table 2: User Requirements for graphics acceleration

Table 3 shows a comparison of the NVIDIA GPU cards to show the relative performance and power of the Quadro K600 through Quadro K6000 devices, and therefore, the typical graphics acceleration requirements for each class of user.

	Quadro K600	Quadro K2000	Quadro K4000	Quadro K5000	Quadro K6000
Stream processors	192	384	768	1536	2880
Texture units	16	32	64	128	240
Raster operation units	16	16	24	32	48
Memory clock	1.8 GHz DDR3	4 GHz GDDR5	5.6 GHz GDDR5	5.4 GHz GDDR5	6 GHz GDDR5
Frame buffer	1 GB	2 GB	3 GB	4 GB	12 GB

Table 3: GPU Performance Comparison

Solution overview

The IBM SDI solution includes the following components:

- User access devices
 - Desktop PCs
 - Thin clients
 - Notebooks
 - Handheld mobile devices
- Virtual infrastructure software
 - Citrix XenDesktop or VMware Horizon View

- Hardware platform
 - IBM iDataPlex server with visualization GPUs
 - IBM NeXtScale System with visualization GPUs

Figure 1 shows the functional components of a SDI solution.

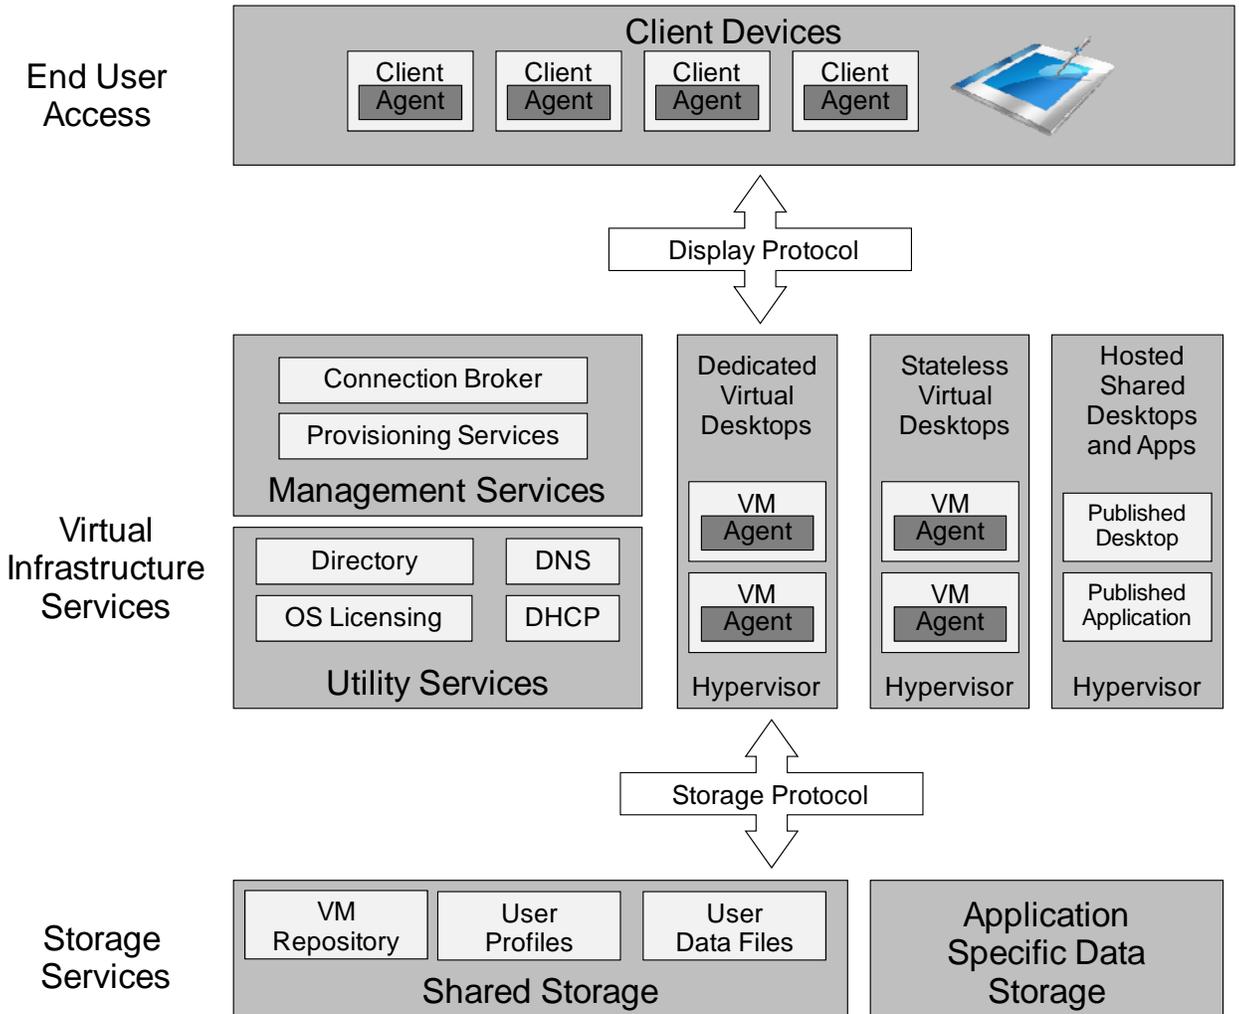


Figure 1: Functional components of the SDI solution

The SDI solution consists of three functional layers:

- **User access layer**
The user access layer is a user entry point into the virtual infrastructure. It includes traditional desktop PCs, thin clients, notebooks, and other handheld mobile devices.
- **Virtual infrastructure services layer**
The virtual infrastructure services layer provides the secure, compliant, and highly available desktop environment to the user. The user access layer interacts with the virtual infrastructure

layer through display protocols such as Remote Desktop Protocol (RDP), PC over IP (PcoIP) and Independent Computing Architecture (ICA).

- **Storage services layer**

The storage services layer stores user persona, profiles, image repositories, and VM disk files. The storage protocol is an interface between the virtual infrastructure services and storage services. The storage protocols include Network File System (NFS), Common Internet File System (CIFS), iSCSI, Fibre Channel, and Fibre Channel over Ethernet (FCoE).

The storage services layer can also provide access to the data that is used by the 3D application. Some applications access the data directly from file or block storage and others must pass through the control of a data manager or database metadata structure. For example, the PTC Creo CAD application makes requests to the PTC Windchill Product Data Management (PDM) system for permission to access and update components.

The virtual infrastructure services layer has the following key functional components:

- **Hypervisor**

The hypervisor provides a virtualized environment for running VMs with the desktop operating systems in them. VMs can either be per user also known as a hosted virtual desktop (HVD) or shared by users and provide either published desktops or published applications.

- **Connection broker**

The connection broker is the point of contact for the client access devices that request the virtual desktops. The connection broker manages the authentication function and ensures that only valid users are allowed access to the infrastructure. When authenticated, it directs the clients to their assigned desktops. If the virtual desktop is unavailable, the connection broker works with the management and provisioning services to have the VM ready and available.

- **Management and provisioning services**

Management and provisioning services allow the centralized management of the virtual infrastructure, providing a single console to manage multiple tasks. These services provide image management, lifecycle management, and monitoring for hosted VMs.

- **High availability services**

High availability (HA) services are part of the connection broker functions and ensure that the VM is up and running even if a critical software or hardware failure occurs. Note that the use of a GPU can make it harder to provide HA because of contextual information around the GPU.

A *dedicated (persistent or personal)* virtual desktop is assigned permanently to the specific user, similar to a traditional desktop PC. Users log in to the same virtual desktop image every time they connect. All changes that they make and each application that they install are saved when the user logs off. The dedicated desktop model is best for users who need the ability to install more applications and store data locally. Users that require substantial 3D graphics acceleration (for example, CAD) are most likely to have dedicated virtual desktops with RAM requirements in the range of 4 to 16 GB depending on the complexity of the application and 3D data.

A *stateless (pooled or non-persistent)* virtual desktop is allocated temporarily to the user. After the user logs off, changes to the image are discarded (reset). Then, the desktop becomes available for the next user, or a new desktop is created for the next user session. A persistent user experience (the ability to personalize the desktop and save data) is achieved through user profile management, folder redirection, and similar approaches. Specific individual applications can be provided to non-persistent desktops by using application virtualization technologies, if required.

Shared virtual desktops or shared virtual applications are a special case of a stateless HVD in which the user shares resources in a single VM with other users. This sharing can be exposed to the user as a desktop or as a single application.

Figure 2 shows the logical connectivity topology that supports the VDI solution.

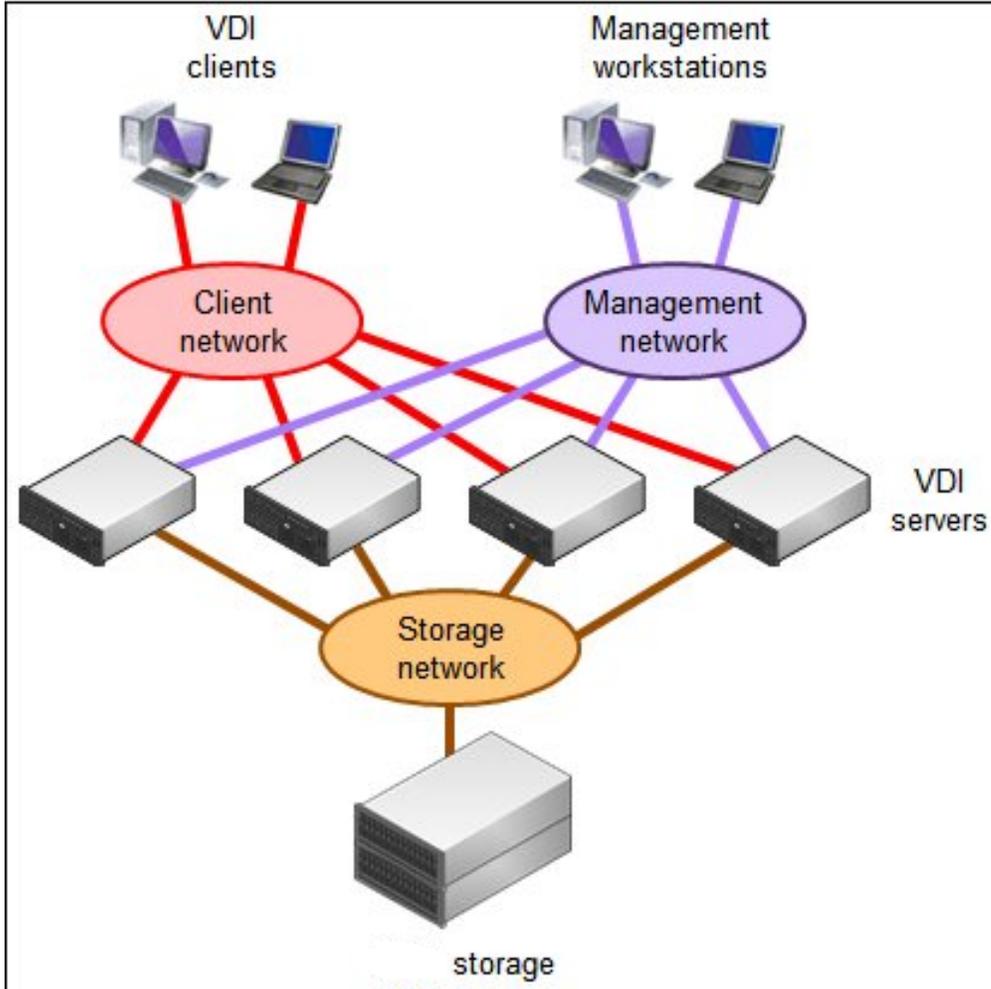


Figure 2: VDI solution logical connectivity topology

The following three networks are defined in the VDI solution.

- Storage network. Carries storage traffic between VDI servers and the NAS system
- Management network. Supports management, provisioning, and monitoring tasks
- Client network. Connects VDI clients (users) to VDI servers

Functional layers and components are supported by a hardware infrastructure platform that must provide the following features:

- Sufficient computing power to support demanding workloads
- Scalability to meet future growth requirements

- Reliability to support business continuity and 24x7 operations
- High-speed, low-latency, and networking for a better user experience
- Cost-efficient storage to handle large amounts of VM and user data
- Centralized management of combined physical and virtual infrastructure from a single user interface to simplify and automate deployment, maintenance, and support tasks

IBM System x iDataPlex or NeXtScale servers, combined with IBM RackSwitch™ networking technologies and IBM storage systems such as the IBM Storwize V7000, represent a robust infrastructure platform that meets these requirements.

Hardware platform

Combined into a single infrastructure, IBM server and networking systems offer unique capabilities that make this infrastructure an exceptional choice for the deployment of the SmartCloud Desktop Infrastructure solution.

Graphics adapters

NVIDIA's Kepler GPU-based GRID K1 and K2 adapters are specifically designed to enable rich graphics in virtualized environments. The K1 and K2 adapters are full-length, double width PCIe x16 adapters.

The GRID adapters feature the NVIDIA Kepler architecture that, for the first time, allows "Hardware virtualization of GPU" – see page 13. This means multiple users can share a single GPU, improving user density while providing true PC performance and compatibility. NVIDIA's patented low-latency remote display technology greatly improves the user experience by reducing the lag that users feel when interacting with their virtual machine. With this technology, the virtual desktop screen is pushed directly to the remoting protocol.

The Kepler GPU includes a high-performance H.264 encoding engine capable of encoding simultaneous streams with superior quality and offloads this work from the processor. The Citrix HDX 3D Pro extensions to ICA use H.264 to reduce data stream bandwidth requirements.

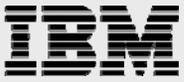
The NVIDIA GRID K1 adapters include four Kepler-based GPUs and 16 GB of DDR3 memory and are designed to host the maximum number of concurrent users. Each of the four K1 GPUs is equivalent to a Quadro K600. The NVIDIA GRID K2 adapters include two higher end Kepler GPUs and 8 GB of DDR5 memory. Each of the two K2 GPUs is equivalent to a Quadro K5000. Refer to Table 3 on page 3 for a comparison of GPU performance.



Figure 3: NVIDIA GRID GPU adapter

Compute servers

IBM System x servers are developed, quality tested, and certified by IBM, a leading vendor in technology patents for 20 years running, and they are backed by IBM's incomparable worldwide service and support.



IBM System x servers deliver business value over the long term because of advancements in scalability, reliability, and performance, combined with flexible configuration options, energy-efficient components, and robust systems management tools.

System x servers provide sufficient processing capacity for the most demanding SmartCloud Desktop Infrastructure deployments requiring graphics acceleration.

IBM iDataPlex dx360 M4

The IBM iDataPlex dx360 M4 server supports the powerful 130 W Intel® Xeon® processor E5-2697 v2, up to 256 GB of memory, and IBM Virtual Fabric I/O connections. The iDataPlex dx360 server supports up to two NVIDIA GRID K1 or K2 GPU adapters.

The PCIe tray attaches to a single iDataPlex dx360 M4 server and forms a 2U server that installs in the 2U chassis. With two 2U riser cards, the combined unit has the following expansions:

- Two double-wide PCIe graphic cards in the upper half of the chassis, slots 3 and 4, each PCIe 3.0 x16 supporting full-height, full-length, double-width PCIe adapters
- Two PCIe cards in the lower half of the chassis, slots 1 and 2, each PCIe 3.0 x8 supporting full-height, half-length PCIe adapters
- Two disk bays on the left for either two 3.5-inch disk drives, four 2.5-inch disk drives, or eight 1.8-inch solid-state drives (SSDs)

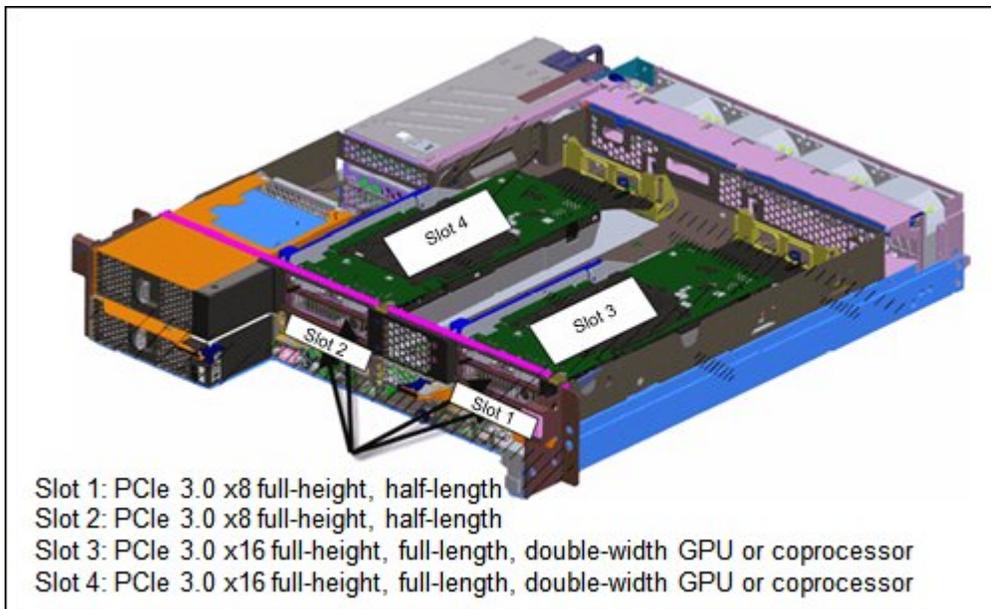
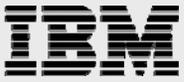


Figure 4: VDI shows the configuration with the PCIe tray

The iDataPlex dx360 M4 2U chassis is designed to be installed in the IBM System x iDataPlex rack. This 100U rack (50U in each half) was specifically designed with a reduced depth to meet high-density data center requirements. It allows infrastructure components, such as switches and power distribution units to be installed into the rack in vertical pockets without sacrificing valuable server space. It is also possible to install the iDataPlex dx360 M4 2U chassis into a standard 19-inch rack.



IBM NeXtScale System

IBM NeXtScale System is a new high-performance computing (HPC) offering from IBM. This new dense offering is based on our experience with IBM iDataPlex and IBM BladeCenter® along with a tight focus on emerging and future client requirements. The IBM NeXtScale n1200 enclosure and IBM NeXtScale nx360 M4 servers are designed to optimize density and performance within typical data center infrastructure limits. The 6U NeXtScale n1200 enclosure fits in a standard 19-inch rack and up to twelve nx360 M4 servers can be installed into the enclosure.

The IBM NeXtScale nx360 also supports the most powerful 130 W Intel Xeon processor E5-2697 v2, up to 256 GB of memory, and IBM Virtual Fabric I/O connections to provide scalable, high-density virtual desktop deployments. The NeXtScale nx360 supports up to two 2.5-inch simple swap SAS/SATA HDDs or SSDs, or four 1.8-inch simple-swap SSDs and up to two NVIDIA GRID K1 or K2 graphics adapters. The two graphics adapters use an extension tray which doubles the height of the NeXtScale nx360. This allows up to six NeXtScale nx360 servers in a 6U NeXtScale n1200 chassis.



Figure 5: NeXtScale nx360 showing GPU tray with two double height GPUs

The six compute nodes are installed in the IBM NeXtScale n1200 enclosure as shown in Figure 6.

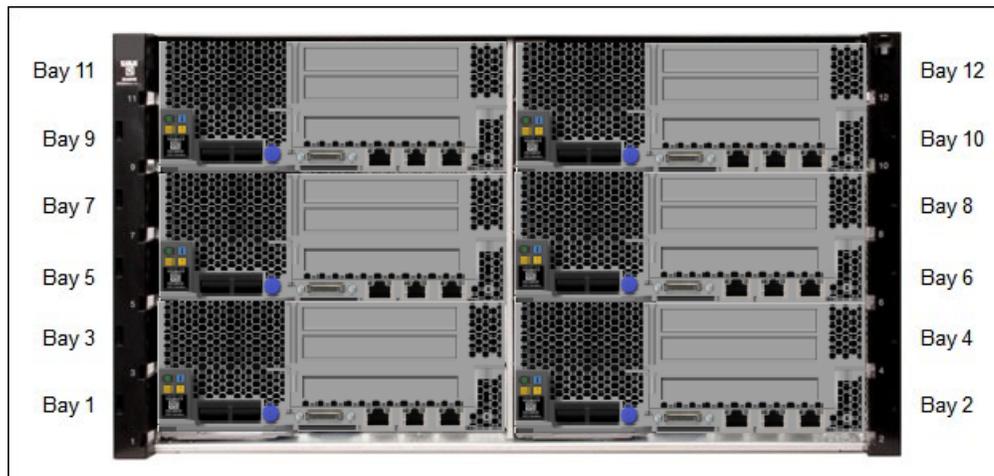


Figure 6: Front view of the IBM NeXtScale n1200 enclosure



Networking

SmartCloud Desktop Infrastructure requires sufficient network bandwidth and efficient traffic management to host as many VMs as possible to ensure that all computing resources are not underused. IBM system networking offerings can help to reduce communication latency and provide the required bandwidth, supporting 10 Gb Ethernet local area network (LAN) connectivity with 10 or 40 Gb uplinks. Virtual Fabric Adapters offer virtual network interface card (NIC) capability to allow up to eight isolated logical ports on a single adapter, with controllable bandwidth allocation per port to manage traffic prioritization.

The bandwidth to the VDI clients is usually much greater for virtual desktops and virtual applications that require graphics acceleration. This is because each *frame* is a rendered image that needs to be sent to the client and each frame could have a very high resolution up to 2560 by 1600 pixels, and more for multiple displays. These users also require a reasonable frame rate per second in order to reduce or eliminate visual artifacts when the graphics image changes. A rate of 24 frames per second is common for the film industry and is sufficient for most users. It is possible to have as few as five frames a second but it becomes hard to use the application at that rate. In general, a 10 GbE LAN connection is recommended for heavy use of VDI graphics acceleration.

Lower speed wide area networks (WANs) can also be used but 2 Mbps per user is a reasonable minimum but for multiple displays and/or higher screen resolutions, the required bandwidth might be upwards of 10 Mbps.

For WANs, it is also important to consider latency. For 24 frames per second, the latency should be 40 milliseconds at most. Most WANs have higher latencies and a 200-millisecond latency would only allow a maximum of 5 frames per second, which is at the very low end of frame speed and application usability. A reasonable objective for most users would be 50 milliseconds; or 20 frames per second.

Shared storage

SDI requires shared storage to store virtual machine images, user home directories, and other desktop related data. The supported file-based protocols are CIFS and NFS with 10 Gb Ethernet and the supported block-based protocols are iSCSI or FCoE with 10 Gb ethernet as well as Fiber Channel. The choice of storage networking is a major decision point for overall system and the type of storage.

The IBM Storwize V3700 and IBM Storwize V7000 storage systems provide block-based shared storage solution that scale from entry virtual desktop environments up to very large enterprise solutions. The Storwize V3700 delivers efficient, entry-level configurations that are designed to meet the needs of small and mid-sized businesses. With the ability to consolidate and share data at an affordable price, Storwize V3700 offers advanced software capabilities that are usually found only in more expensive systems. The Storwize V7000 system delivers enterprise-level storage with up to 960 hard disk or flash drives. Both Storwize systems offer a large flash cache and IBM Easy Tier® which can store the most frequently accessed data on flash drives. Windows Storage Server 2012 can be deployed as a failover cluster with two IBM servers and the external shared Storwize V3700 or Storwize V7000 storage system. The high availability cluster performs file serving tasks hosting CIFS and NFS file shares for Windows-based VDI users.

The IBM System Storage N series systems provide powerful virtualization and thin provisioning capabilities to help you maximize storage utilization and minimize the use of power, cooling, and floor space. At the same time, you can improve staff productivity with an integrated suite of application-aware manageability software that offers policy-based automation to otherwise manual tasks, improving storage efficiency.



Solution architecture

IBM SDI with Citrix XenDesktop or VMware Horizon View can help to transform Microsoft Windows desktops and applications that require graphics acceleration into a cloud-type service that is accessible on virtually any device, anywhere.

Both Citrix XenDesktop and VMware Horizon View provide a method to use software emulation of a GPU. This can be processor-intensive and disruptive to other users who have a choppy experience because of reduced processor performance. Software emulation is not recommended for any user who requires graphics.

Virtual Desktop Graphics Acceleration

For users who need virtual desktops with graphics acceleration, there are three options of how to use the GPU as summarized in Table 4. The best option for a particular application, user, or environment depends upon its performance characteristics and there is no single prescriptive choice.

Characteristic	Dedicated GPU	Shared GPU	GPU hardware virtualization
Performance/cost per user	Best performance	Low cost per user	Flexible choice of cost per user and performance.
Guaranteed GPU performance	Yes	No	Yes
Supports motion (live migration of VMs)	No	Yes	No
All versions of 3D APIs supported (DirectX and OpenGL)	All	Limited	All
Recommended for Designer/Engineer	Yes	No	Yes
Recommended for GPU power user	No	No	Yes
Recommended for GPU knowledge worker	No	No	Yes
ESXi hypervisor support	Yes	Yes	No
XenServer hypervisor support	Yes	No	Yes
Citrix XenDesktop/XenApp support	XenDesktop	XenDesktop	XenDesktop
VMware Horizon View support	Yes	Yes	No

Table 4: Summary of options for virtual desktop graphical acceleration

Dedicated GPU

With desktop PCs and workstations, a user has complete and sole access to a GPU. In a virtualized environment, the same pattern can be used where a user has access to their own dedicated GPU. Naturally, this limits the number of users on a server to the number of GPUs available on that server. For example an iDataPlex dx360 M4 server has a capacity of two GRID K2 adapters and can therefore support up to four users per server.

The dedicated GPU option as depicted in Figure 7 requires support from the hypervisor and is most often used by designers, engineers, or scientists who need the best possible performance and work in fields such as CAD, financial trading, CAE, petrochemical, and other technical computing applications.

The VMware ESXi hypervisor 5.x implements a virtual dedicated graphics acceleration (vDGA) mode that can be used by both VMware Horizon View and Citrix XenDesktop. The Citrix XenServer 6.x hypervisor implements a GPU pass-thru mode that can only be used by Citrix XenDesktop.

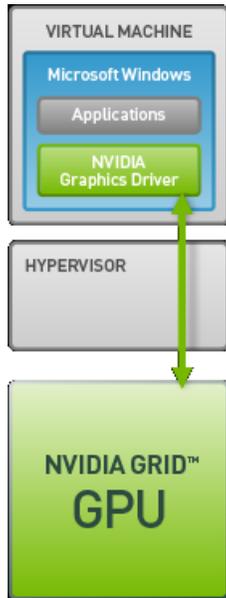


Figure 7: Dedicated GPU

Shared GPU

VMware ESXi hypervisor 5.1 and later also supports a GPU sharing mode called virtual shared graphics acceleration (vSGA) as shown in Figure 8. vSGA is a form of software virtualization of a GPU and can be used with both VMware View Horizon and Citrix XenDesktop. The number of users is limited by the processor, memory, and GPU usage in the server.

At any particular time, multiple users might be trying to access the GPU at the same time. This means that there can be choppy performance because each user is essentially vying for dedicated use of the GPU. For this reason, the vSGA GPU sharing option is not generally recommended.

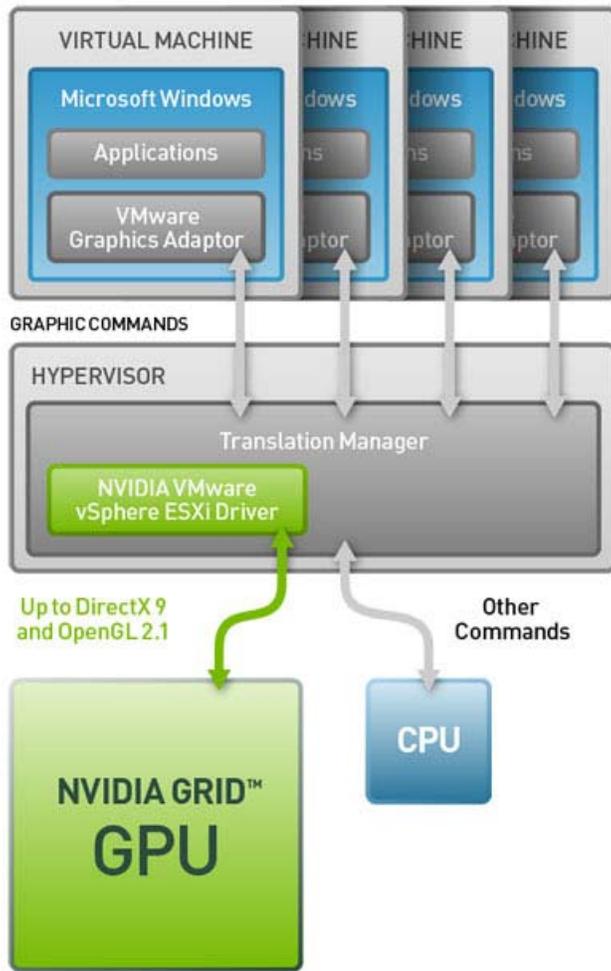


Figure 8: VMware ESXi vSGA (shared GPU)

Hardware virtualization of GPU

Software virtualization of a GPU can be problematic because of contention for GPU resources by multiple users. This problem is resolved by using hardware virtualization whereby the hardware is partitioned and each user only has access to their partition. The NVIDIA GRID GPUs were specifically designed to manage this partitioning and ensure each user can have simultaneous, unimpeded access to their assigned partition. The NVIDIA vGPU manager can assign up to eight users per GPU chip.

NVIDIA has defined profiles that specify the size of each partition, and therefore, the number of simultaneous users for each GRID adapter. Because the number of users for a given profile is fixed, the maximum number of users is also fixed. Table 5 gives the name of the profile, the type of GRID adapter (K1 or K2), the number of users per adapter and the allocated GPU resources. The last column gives an indication of the type of user for that profile.

GRID adapter	Profile name	Number of users	Graphics memory	Maximum displays	Maximum resolution	Use cases
K2	K260Q	4	2 GB	4	2560 x 1600	Designer/Engineer
	K240Q	8	1 GB	2	2560 x 1600	GPU power user
	K220Q	16	0.5 GB	2	2560 x 1600	GPU power user
	K200	16	0.25 GB	2	1900 x 1200	GPU knowledge worker
K1	K140Q	16	1 GB	2	2560 x 1600	GPU power user
	K120Q	32	0.5 GB	2	2560 x 1600	GPU power user
	K100	32	0.25 GB	2	1900 x 1200	GPU knowledge worker

Table 5: Virtual GPU profiles

The profiles ending with the letter “Q” require certification of applications. These users were probably already using a graphics workstation with an included GPU. The GPU usage requirements determine which profile to use.

For designers/engineers, it is recommended that the K260Q profile or a dedicated GPU should be used. A server with two GRID K2 adapters can support eight users, with two users sharing each physical GPU.

For GPU power users, it is recommended that the K240Q, K220Q, K140Q, or K120Q profile is used. A server with two GRID K1 adapters can support up to 64 users, with eight users sharing each physical GPU.

For GPU knowledge workers, it is recommended to use the K100 or K200 profile. A server with two GRID K1 adapters can support up to 64 users, with eight users sharing each physical GPU.

Host Shared Graphics Acceleration

Citrix, using its XenApp technology, also supports hosted shared desktops and hosted shared applications whereby one or more users can share a virtual machine. The VM may utilize a dedicated GPU using either ESXi with vDGA, or XenServer with pass-thru mode as shown in Figure 9 below. In this case, more than one user has access to the VM and shares the GPU with other users.



Figure 9: Citrix XenApp Sharing of a GPU

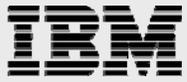


Table 6 summarizes the characteristics of using XenApp with a dedicated GPU.

Characteristic	XenApp use of a dedicated GPU
Performance/cost per user	Good
Guaranteed GPU performance	Yes
Supports motion (live migration of VMs)	No
All versions of 3D APIs supported (DirectX and OpenGL)	All
Recommended for Designer/Engineer	No
Recommended for GPU power user	No
Recommended for GPU knowledge worker	Yes
ESXi hypervisor support	Yes
XenServer hypervisor support	Yes
Citrix XenDesktop/XenApp support	XenApp
VMware Horizon View support	No

Table 6: Summary for hosted shared virtual desktop/virtual application graphical acceleration

Citrix XenDesktop

By using the open architecture of Citrix XenDesktop, customers can adopt desktop virtualization quickly and easily with any hypervisor, storage, or management infrastructure.

The following XenDesktop features provide a familiar experience for the user:

- Multiple monitor support
- 3D graphics business application support
- Multimedia support
- Printing from a virtual desktop
- Accessing USB devices and other peripheral devices such as a 3D mouse
- Roaming user profiles
- XenDesktop offers several levels of security features, including:
- Multifactor authentication
- Traffic encryption
- Built-in password management
- Secure Sockets Layer (SSL) tunneling to ensure that all connections are encrypted

The following Citrix XenDesktop features provide centralized administration and management:

- Microsoft Active Directory
- Web-based administrative console
- Automated desktop provisioning and storage optimization

XenDesktop includes the following scalability, integration, and optimization features:

- VMware vSphere, Microsoft Hyper-V, and XenServer hypervisor support
- Integration with VMware vCenter to achieve cost-effective densities, high levels of availability, and advanced resource allocation control for virtual desktops
- Automated provisioning of desktop images that share virtual disks with a master image

For graphics acceleration, in particular, XenDesktop offers a range of solutions including support for dedicated GPU and software virtualized GPU. Citrix XenDesktop also offers hardware virtualized GPU support with XenServer 6.2 onwards.

VMware Horizon View

The following VMware Horizon View features provide a familiar experience for the user:

- Take advantage of multiple monitor support for RDP and PCoIP, where with PCoIP, you can adjust the display resolution and rotation separately for each monitor.
- Print from a virtual desktop (in a Microsoft Windows environment) to any local or networked printer.
- Access USB devices and other peripheral devices that are connected to the local device that displays your virtual desktop.
- Manage profiles by using View Persona Management to preserve user profiles and data between sessions and to dynamically synchronize them to a remote CIFS share at configurable intervals. View Persona Management can work with or without Windows roaming profiles.

VMware View offers several levels of security features, including the following features:

- Two-factor authentication, such as RSA SecurID, RADIUS, or smart cards
- Precreated Active Directory accounts to provision View desktops in environments that have read-only access policies for Active Directory
- SSL tunneling to ensure that all connections are encrypted

The following VMware View features provide centralized administration and management:

- Microsoft Active Directory
- Web-based administrative console
- Use of a template or master image to quickly create and provision pools of desktops, virtual desktop updates, and patches

Scalability features depend on the VMware virtualization platform to manage both desktops and servers:

- You can integrate with VMware vCenter to achieve cost-effective densities, high levels of availability, and advanced resource allocation control for your virtual desktops.
- You can use View Composer to quickly create desktop images that share virtual disks with a master image. By using linked clones in this way, you conserve disk space and simplify the management of patches and updates to the operating system.

For graphics acceleration in particular, VMware Horizon View offers software virtualized GPUs (vSGA mode) and dedicated GPUs (vDGA mode).

Reference design

The configurations described in this solution guide provide a balanced design point that does not overstress the processor, memory, storage, or GPU. If users require larger desktops, then more system memory can be added to each server. Paging should be avoided as it affects overall performance. When hardware virtualization is used, it might be possible to choose an alternative profile to give better performance or GPU utilization.

High availability of a solution is required for business continuity. High availability needs to be built end to end and some parts are outside the scope of this solution design.

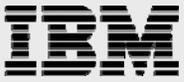
Depending on the application and user, there is a wide range of performance requirements for graphics acceleration, and therefore, the number and type of GPU adapters needed for a particular solution. The following guidelines can be used to best determine the GPU adapter:

- For designers/engineers/scientists the predominant option is expected to be dedicated access to a GPU. For IBM servers that support two NVIDIA GRID K2 adapters, this is a maximum of four users per server.
- For users who have casual use for a GPU, then it is possible to use software virtualization of the GPU. The number of users in this case depends more on the processor and memory in the server.
- If the power user already has a workstation with a GPU, then hardware virtualization of the NVIDIA GRID K2 can be used. For IBM servers that support two NVIDIA GRID K2 adapters, this is between 4 and 32 users per server. The choice of a virtual GPU profile is dependent on the GPU performance requirements.
- For knowledge workers with applications that would benefit from a GPU such as multimedia content providers, then hardware virtualization of the NVIDIA GRID K1 can be used. For IBM servers that support two NVIDIA GRID K1 adapters, this is between 32 and 64 users per server. Note that in this case, the maximum number of users can be driven more by limitations on the processor compute power. For example, if each user needs three vCPUs, then two 12 core processors can support at most 16 users and one NVIDIA GRID K1 is sufficient.

The amount of memory for these solutions is dependent on the number of users and the memory in each VM. Although these solutions have fewer users than non-GPU solutions, the amount of memory in each VM can be much greater.

The following guidelines can be used to best determine the server memory:

- Designers/Engineers might need between 8 GB and 16 GB of memory. However, there are only four users per server. In this case, 64 GB or 128 GB of memory is sufficient.
- For software virtualization of the GPU using XenApp, the required server memory is 128 GB. This is because the hosted shared desktops or applications are bounded by the processor's performance rather than memory.
- For software virtualization of the GPU using vSGA, the number of users is also limited by processor performance. In this case, the server might require as much as 384 GB of memory to support 95 users with 4 GB VMs.
- In general, users with greater GPU performance demands also require more memory in the VM. Users of the GRID K2 might require between 4 GB and 8 GB of memory. With up to 32 users per server, the maximum memory requirement is 256 GB. Users of the GRID K1 might require between 2 GB and 4 GB of memory. With up to 64 users per server, the maximum memory requirement is also 256 GB.



A proof of concept with a smaller number of servers can be used to verify the configuration. An initial estimate can be derived from the GPU and memory rules of thumb listed earlier.

It is assumed that Active Directory servers already exist. If not, then a fully redundant pair of Active Directory servers must be created.

The following section provides two examples of reference designs. One is an aircraft company and other is a publishing company. Note that the networking, shared storage, and application storage are all solution dependent.

Aircraft company example

An aircraft company has a large team of CAD designers. The cockpit design is being done by a partner in a remote location. The aircraft company decides to use remote virtual desktops for the cockpit design company so their 20 designers can access the total design. Note that the design files stay within the data center of the aircraft company thus reducing the risk of intellectual property theft.

In this case, it has been determined to use dedicated GPUs and the NVIDIA GRID K2 to ensure the best performance for the CAD application. The cockpit designers in the remote location have a dedicated fast connection into the aircraft company data center.

The aircraft company already has an iDataPlex cluster that is used for CAE computation of fluid flows and airframe stresses. They want to simply add some more iDataPlex nodes for the CAD 3D visualization. There is also plenty of storage capacity and I/O on their existing system.

The reference design for this example consists of:

- One iDataPlex chassis with two 1U planars for fault-tolerant VDI management functions
- Six iDataPlex chassis with base planar and PCIe expansion tray for two GRID K2 adapters. One extra compute nodes is for failover purposes in case of a problem.
- VMware ESXi 5.5 hypervisor and VMware Horizon View 5.3 using vDGA mode

Reference	Description
①	One iDataPlex chassis with two 1U planars
	Two Intel Xeon processor E5-2650 v2 Eight 16 GB memory DIMMs Two integrated 10 Gb converged network ports One embedded ESXi hypervisor on a USB key
②	Six iDataPlex chassis with base planar with two GRID K2 adapters
	Two Intel Xeon processor E5-2680 v2 Eight 16 GB memory DIMMs Two integrated 10 Gb converged network ports One embedded ESXi hypervisor on a USB key One PCI expansion tray Two GRID K2 adapters



Publishing company example

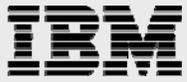
A publishing company would like all of their 250 digital content editors and magazine editors to use a VDI solution so that they can work from home or in the field using a laptop or tablet. All users are required to use stateless (non-persistent) virtual desktops for easier management and each virtual desktop is 4GB.

In this case it has been determined to use hardware virtualized GPUs and the NVIDIA GRID K1. The K140Q profile with up to 32 users per server provides the best performance. In order to allow for failover, 10 compute nodes are used for a capacity of up to 320 users. Up to two of these compute nodes can fail and that still leaves 8 which can service up to 256 users (8 x 32).

The reference design for this example consists of the following:

- Two NeXtScale n1200 chassis
- Two single-height NeXtScale nx360 M4 compute nodes for fault-tolerant VDI management functions
- Two single-height NeXtScale nx360 M4 compute nodes for fault-tolerant Windows Storage Server 2012 R2
- Ten double-height NeXtScale nx360 M4 compute nodes with two GRID K1 adapters. Two extra compute nodes are used for failover purposes in case of a problem.
- Two G8124E 10 GbE Ethernet top of rack switches
- One IBM Storwize V3700 storage system with two expansion enclosures
- Citrix XenServer 6.2 hypervisor and Citrix XenDesktop 7.1 using hardware virtualization

Reference	Description
①	Two NeXtScale n1200 chassis
	Six 1300W power supplies and power cables Ten fan assemblies KVM dongle cable
②	Two NeXtScale nx360 M4 for VDI management functions
	Two Intel Xeon processor E5-2650 v2 Eight 16 GB memory DIMMs Two integrated 10 Gb Ethernet ports with LOM upgrade Two 300 GB 10,000 rpm SAS HDDs for the hypervisor
③	Two NeXtScale nx360 M4 for Windows Storage Server 2012 R2
	Two Intel Xeon processor E5-2650 v2 Four 8 GB memory DIMMs Two integrated 10 Gb Ethernet ports with LOM upgrade



	Two 300 GB 10,000 rpm SAS HDDs for the hypervisor
④	Ten NeXtScale nx360 M4 with two GRID K1 adapters Two Intel Xeon processor E5-2680 v2 Eight 16 GB memory DIMMs Two integrated 10 Gb Ethernet ports with LOM upgrade Two 300 GB 10,000 rpm SAS HDDs for the hypervisor One PCIe Native Expansion (NeX) Tray Two GRID K1 adapters
⑤	Two 10 GbE Ethernet top of rack switches 10 GbE cabling to NeXtScale nx360 nodes and Storwize V3700 storage
⑥	One IBM Storwize V3700 storage system with one expansion enclosure Four 10 Gb iSCSI ports per controller Sixty-four 300 GB 10K rpm SAS 2.5-inch HDDs in a RAID-10 Eight 400GB SSDs for EasyTier



Ordering information

Aircraft company example

The following table shows the part numbers and quantities for ordering the example configuration for the aircraft company.

Reference	Part number	Description	Quantity
①	One iDataPlex chassis with two 1U planars		
	7913HC1	IBM System x iDataPlex dx360 M4 2U chassis	1
	A1G7	IBM System x 900W High Efficiency Platinum AC Power Supply	2
	A1V4	AC Inlet Y 12A Cable	2
	A262	iDataPlex 2U Dual Planar Configuration	1
	7912AC1	IBM System x iDataPlex dx360 M4 server	2
	A3QL	16GB (1x16GB, 2Rx4, 1.5V) PC3-14900 CL13 1866MHz LP RDIMM	16
	A40M	Intel Xeon Processor E5-2650 v2 8C 2.6GHz 20MB Cache 1866MHz 95W	4
	A1ML	IBM Integrated Management Module Advanced Upgrade	2
	A4BV	IBM System x iDataPlex dx360 M4 (Base 3)	2
	A1UM	2.5" HDD Cage Assembly	2
	5977	Select Storage devices - no IBM-configured RAID required	2
	A584	IBM USB Memory Key for VMWare ESXi 5.5	2
	5751	QLogic 10Gb CNA for IBM System x	2
	A1T3	IBM System x iDataPlex dx360 M4 server	2
	A262	iDataPlex 2U Dual Planar Configuration	2
	A1UD	Planar Tray Top Cover Assembly	2
A1UE	IDPX PCI Riser Module	4	
②	Six iDataPlex chassis with base planar and two GRID K2 adapters		
	7913HC1	IBM System x iDataPlex dx360 M4 2U chassis	6
	A1G7	IBM System x 900W High Efficiency Platinum AC Power Supply	12
	A1V4	AC Inlet Y 12A Cable	12
	A2N3	iDataPlex 2U GPU Configuration	6
	A1UW	Power Rating Label - GPU Tray (900W)	6
	7912AC1	IBM System x iDataPlex dx360 M4 server	6
	A3QL	16GB (1x16GB, 2Rx4, 1.5V) PC3-14900 CL13 1866MHz LP RDIMM	48
	A3ZC	Intel Xeon Processor E5-2680 v2 10C 2.8GHz 25MB Cache 1866MHz 115W	12



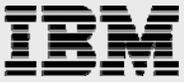
Reference	Part number	Description	Quantity
	A1ML	IBM Integrated Management Module Advanced Upgrade	6
	A4BV	IBM System x iDataPlex dx360 M4 (Base 3)	6
	A1UM	2.5" HDD Cage Assembly	6
	5977	Select Storage devices - no IBM-configured RAID required	6
	A584	IBM USB Memory Key for VMWare ESXi 5.5	6
	5751	QLogic 10Gb CNA for IBM System x	6
	A1UG	IDPX PCI Tray	6
	A2N3	iDataPlex 2U GPU Configuration	6
	A3GN	NVIDIA GRID K2	12
	A1UH	iDPx 2U GPU PCI Riser	12

Publishing company example

The following table shows the part numbers and quantities for ordering the example configuration for the publishing company.

Reference	Part number	Description	Quantity
①	Two NeXtScale n1200 chassis		
	5456HC1	IBM n1200 Enclosure Chassis Base Model	2
	A4MM	CFF 1300W Power Supply	12
	6201	1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	12
	A41D	n1200 Enclosure Chassis	2
	2306	Rack Installation >1U Component	2
	A4AK	KVM Dongle Cable	2
	A41F	n1200 Enclosure Fan Assembly	20
	A41H	n1200 Enclosure Midplane Assembly	2
	A41G	n1200 Enclosure fan Power Control Card Assembly	2
②	Two nx360 M4 for VDI management functions		
	5455AC1	nx360 M4 Computer Node	2
	A3QL	16GB (1x16GB, 2Rx4, 1.5V) PC3-14900 CL13 1866MHz LP RDIMM	16
	5978	-SB- Select Storage devices - IBM-configured RAID	2
	8947	-SB- Integrated SATA Single Mirroring - 2 identical SS HDDs required	2
	A48B	-SB- IBM 500GB 7.2K 6Gbps SATA 2.5" HDD for NeXtScale System	4
	A426	Intel Xeon Processor E5-2650 v2 8C 2.6GHz 20MB Cache 1866MHz 95W	2

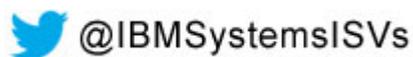
Reference	Part number	Description	Quantity
	A42G	Addl Intel Xeon Processor E5-2650 v2 8C 2.6GHz 20MB 1866MHz 95W	2
	A41J	xc360 M4 3.5-inch HDD cage	2
	A17T	ServeRAID C100 for System x	2
	A1MK	IBM Integration Management Module Standard Upgrade	2
	A1ML	-SB- IBM Integrated Management Module Advanced Upgrade	2
	A41E	nx360 M4 Computer Node	2
	A22H	Qlogic Dual Port 10GbE SFP+ Embedded VFA for IBM System x	2
	A2TF	Qlogic Embedded VFA FCoE/iSCSI License for IBM System x (FoD)	2
3	Two nx360 M4 for Windows Storage Server 2012 R2		
	5455AC1	nx360 M4 Computer Node	2
	A3QJ	8GB (1x8GB, 2Rx8, 1.5V) PC3-14900 CL13 1866MHz LP RDIMM	10
	5978	-SB- Select Storage devices - IBM-configured RAID	2
	8947	-SB- Integrated SATA Single Mirroring - 2 identical SS HDDs required	2
	A48B	-SB- IBM 500GB 7.2K 6Gbps SATA 2.5" HDD for NeXtScale System	4
	A426	Intel Xeon Processor E5-2650 v2 8C 2.6GHz 20MB Cache 1866MHz 95W	2
	A42G	Addl Intel Xeon Processor E5-2650 v2 8C 2.6GHz 20MB 1866MHz 95W	2
	A41K	xc360 M4 2.5-inch HDD cage	2
	A17T	ServeRAID C100 for System x	2
	A1MK	IBM Integration Management Module Standard Upgrade	2
	A1ML	-SB- IBM Integrated Management Module Advanced Upgrade	2
	A41E	nx360 M4 Computer Node	2
	A22H	Qlogic Dual Port 10GbE SFP+ Embedded VFA for IBM System x	2
	A2TF	Qlogic Embedded VFA FCoE/iSCSI License for IBM System x (FoD)	2
4	Ten nx360 M4 with two GRID K1 adapters		
	5455AC1	nx360 M4 Computer Node	10
	A3QL	16GB (1x16GB, 2Rx4, 1.5V) PC3-14900 CL13 1866MHz LP RDIMM	80
	5978	-SB- Select Storage devices - IBM-configured RAID	10
	8947	-SB- Integrated SATA Single Mirroring - 2 identical SS HDDs required	10
	A48B	-SB- IBM 500GB 7.2K 6Gbps SATA 2.5" HDD for NeXtScale System	20
	A429	Intel Xeon Processor E5-2680 v2 10C 2.8GHz 25MB Cache 1866MHz 115W	10
	A42K	Addl Intel Xeon Processor E5-2680 v2 10C 2.8GHz 25MB 1866MHz 115W	10
	A41K	xc360 M4 2.5-inch HDD cage	10

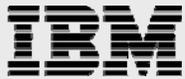


Reference	Part number	Description	Quantity
	A17T	ServeRAID C100 for System x	10
	A1MK	IBM Integration Management Module Standard Upgrade	10
	A1ML	-SB- IBM Integrated Management Module Advanced Upgrade	10
	A41E	nx360 M4 Computer Node	10
	A41R	nx360 M4 PCIe riser	10
	A22H	Qlogic Dual Port 10GbE SFP+ Embedded VFA for IBM System x	10
	A2TF	Qlogic Embedded VFA FCoE/iSCSI License for IBM System x (FoD)	10
	A4MB	IBM NeXtScale PCIe Native Expansion Tray	10
	A3GM	NVIDIA GRID K1	20
5	Two 10 GbE Ethernet top of rack switches		
	4668HCF	Edge-Core ECS5610-52S 10GbE Switch (PSE)	1
	6311	2.8m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	1
	3737	3m Molex Direct Attach Copper SFP+ Cable	36
	49Y7890	1m IBM QSFP+-to-QSFP+ cable	1
6	One IBM Storwize V3700 storage system with one expansion enclosure		
	207224C	IBM Storwize V3700 SFF Dual Control Enclosure	1
	ACHB	4GB to 8GB Cache Upgrade	2
	ACLB	300GB 15K 2.5 Inch HDD	16
	ACMB	400GB 2.5 inch SSD	8
	ACHM	10Gb iSCSI - FCoE 2 Port Host Interface Card	2
	ACFN	-SB- Easy Tier	1
	207224C	IBM Storwize V3700 SFF Dual Control Enclosure	1
	ACHB	4GB to 8GB Cache Upgrade	2
	207224E	IBM Storwize V3700 SFF Expansion Enclosure	2
	ACLB	300GB 15K 2.5 Inch HDD	48

Resources

- [IBM SmartCloud Desktop Infrastructure reference architecture](#)
ibm.com/partnerworld/page/stg_ast_eis_sdi_infrastructure
- [IBM iDataPlex servers](#)
ibm.com/redbooks/abstracts/tips0878.html
- [IBM NeXtScale nx360 M4 servers](#)
ibm.com/redbooks/abstracts/tips1051.html
- [NVIDIA GRID K1 and K2 adapters](#)
nvidia.com/object/grid-boards.html#sthash.Dq0AMR7W.dpuf
- [NVIDIA GRID Certifications](#)
nvidia.com/content/cloud-computing/pdf/GRID_Certifications.pdf
- [Citrix XenDesktop](#)
citrix.com/products/xendesktop
- [Citrix graphics acceleration white paper](#)
citrix.com/content/dam/citrix/en_us/documents/products-solutions/virtualize-3d-professional-graphics-design-guide.pdf
- [VMware vSphere](#)
vmware.com/products/datacenter-virtualization/vsphere
- [VMware Horizon View](#)
vmware.com/products/horizon-view
- [VMware graphics acceleration white paper](#)
vmware.com/files/pdf/techpaper/vmware-horizon-view-graphics-acceleration-deployment.pdf
- [IBM System Storage Interoperation Center \(SSIC\)](#)
ibm.com/systems/support/storage/ssic/interoperability.wss





Trademarks and special notices

© Copyright IBM Corporation 2014.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

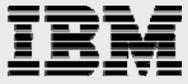
Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.



Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.