

ALTERNATE CLIENT ARCHITECTURES OVERVIEW

DELIVERING APPLICATIONS TO THE DESKTOP

Enterprise IT departments are tasked with delivering applications, data, and services to users in a reliable, secure fashion. As the enterprise has evolved, the method for delivery has too. In today's enterprise, myriad architectural options exist for IT managers and implementers to choose from. These technologies include operating system streaming, application streaming, virtualization and isolation, Virtual Desktop Infrastructure, Terminal Server-based computing, operating system isolation, Web-based services, and many, many others. Each technology has its strengths and its challenges; and each can provide performance and cost benefits in certain situations.

This white paper focuses on four of the prevalent alternative architectures: Virtual Desktop Infrastructure (VDI), Terminal Services, operating system (OS) streaming, and Blade Personal Computers (PCs). This white paper presents a technical overview and potential use cases for each of these architectures, and explores AMD as the platform of choice for each by demonstrating both technical (performance) and non-technical (cost) advantages.

INTRODUCTION

Many organizations struggle with the day-to-day management of user desktops. Traditional IT tasks, such as procurement of hardware, distribution and patching of software, inventory management, updating antivirus signatures, and supporting desktops, have become more complex with the added challenges of securing the network, desktop, and data. As companies grow and user populations disperse, these problems can grow in an exponential fashion. This problem manifests itself in other areas as well. Consider the company that has hired on a consultant to develop a customized application. Providing that consultant with a powerful work environment in a remote setting—all while securing the parent organization's source code and intellectual property—can be a daunting challenge.

Businesses have been exploring a wide variety of ways to address this problem. The result has been a multitude of client computing architectures. This white paper surveys four popular architectures to help IT administrators navigate their myriad of architectural options. This paper seeks to clarify these architectural paradigms, describe who the players are, explain what the design options are, and explore which circumstances best lend themselves to these various alternatives. This paper also shows how AMD, a leading manufacturer of server processors, supports and enhances these architecture alternatives. After reading this paper, IT administrators and decision makers should have a better framework for evaluating and selecting the client computing architectures that are best suited for their needs.

The four architectures described in this white paper—Virtual Desktop Infrastructure (VDI), Terminal Services, operating system (OS) streaming, and Blade Personal Computers (PCs)—provide alternatives to the traditional commercial desktop and address the issues of scalability, security, reliability, and cost in very different ways. All four provide the full desktop experience to users, but without the need for a locally executed copy of the operating system.

TABLE OF CONTENTS

Delivering Applications to the Desktop	1
Introduction	1
AMD and the Role of Hardware	3
Virtual Desktop Infrastructure	3
What Is VDI?	3
VDI Use Cases	4
VDI Advantages	4
VDI Disadvantages	5
VDI Providers	5
AMD Value Proposition	5
Terminal Services	6
What Is Terminal Services?	6
Terminal Services Use Cases	7
Terminal Services Advantages	7
Terminal Services Disadvantages	7
Terminal Services Providers	7
Operating System Streaming	8
What Is Operating System Streaming?	8
Operating System Streaming Use Cases	9
Operating System Streaming Advantages	9
Operating System Streaming Disadvantages	9
Operating System Streaming Providers	9
Blade PC	10
What Are Blade PCs?	10
Blade PC Use Cases	11
Blade PCs Advantages	11
Blade PC Disadvantages	11
Blade PC Providers	11
Client Technology Considerations	11
Amd Technology Advantages	12
Summary	13
Glossary	16

AMD AND THE ROLE OF HARDWARE

The alternate architectures described in this white paper are solutions that can be further enhanced with an appropriate underlying platform. At the core of enterprise computing, regardless of the architecture, is the server and the server processor. AMD manufactures processors with a wide range of technology options that can be tailored to enhance a variety of computing architectures. AMD processors support the performance and isolation requirements of the virtualization technology used in VDI and Terminal Services, in addition to the low power and memory requirements for concurrent users and applications demanded by Blade PCs and OS streaming.

The AMD Opteron™ processor family of server-class processors based on the 64-bit architecture is compatible with pre-existing x86 32-bit architecture. AMD Opteron processors utilize Direct Connect Architecture, which helps to eliminate the challenges and bottlenecks associated with traditional front-side bus system architectures. The AMD Opteron processor introduces several new features, including quad-core upgradeability, AMD Virtualization™ (AMD-V™) technology, and energy-efficient double-data-rate two (DDR2) memory. In addition, AMD Opteron processors are designed to advance performance-per-watt (PPW) capabilities and apply the proven technologies introduced with earlier AMD Opteron processors. AMD Opteron processors are a preferred platform for desktop delivery solutions, providing cost and performance advantages regardless of the architecture solution.

VIRTUAL DESKTOP INFRASTRUCTURE

The role of an IT department is to provide applications and services to users. In order to run an application, a user needs a desktop—be it a Windows® desktop, a browser window, or some other backdrop—that lets the user select and launch the desired applications. Virtual Desktop Infrastructure (VDI) delivers a full, rich desktop experience that is hosted on a centralized data center server, helping increase security and decrease total cost of ownership (TCO).

WHAT IS VDI?

In the VDI model, a hardware virtualization layer is added to the data center server (or servers) in place of a more traditional operating system. This virtualization layer provides numerous virtual machines (VMs), each with an operating system, applications, device

configurations, and a unique desktop environment (or graphical user interface [GUI], including customizations such as Windows wallpaper and screen savers) for each user (see Figure 1). Because the VM is a complete instance of a PC, the user session functions just as a locally run version of the operating system would.

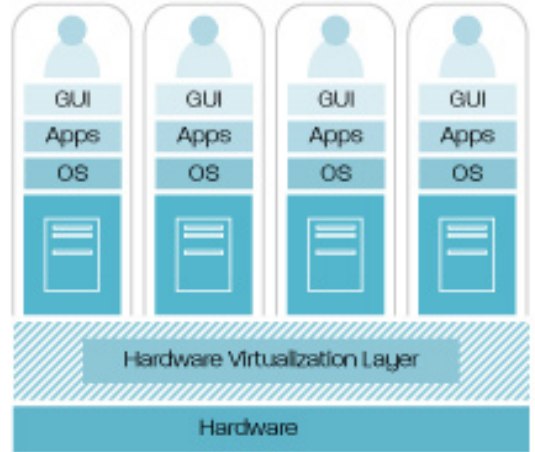


Figure 1 VDI structure

Users access their virtual desktops remotely from a traditional PC, thin client, or repurposed computer through a connection broker, which connects the user to the appropriate VM. The virtual desktop is delivered to the client either using Remote Display Protocol (RDP), or directly from the connection broker (as in the case of Citrix Presentation server, which is built on Independent Computing Architecture [ICA], a thin client protocol). The hosted desktop images are completely isolated from each other on the secure data center server; each hosted VM runs independently on the server to protect against unauthorized access or system crashes (see Figure 2).

Virtual Desktop Infrastructure

- Each user has his or her own operating system instance
- Applications execute on each user's VM




Figure 2 VDI architecture

VDI USE CASES

With VDI, a user can run any earlier, off-the-shelf, or custom application without modification. VDI works well for unique desktops or applications that have unusual system requirements; earlier applications that would need to be rewritten or that simply do not run in other server-based computing (SBC) environments can be run as-is with VDI.

VDI can enhance a company's disaster recovery (DR) plan. In the case of a disaster, be it natural or electronic, secure and continued access to the organization's desktop environments is critical to reduce disruption. With VDI, hosted desktops in VMs can be easily moved or redeployed to ensure quick recovery and business continuity.

Today's highly mobile workforce is placing increasing demands for corporate data access from remote locations. VDI enables employees using mobile devices to securely access their desktops, along with their authorized applications and associated data, without the risk of compromising corporate assets if their devices are lost or stolen.

VDI bridges the gap between the users' desire to have a personalized, isolated desktop experience and the IT administrators' need to maintain a secure, centrally controlled and managed computing environment. Outsourced or offshore activities (call center operations, order and back-office processing, or software development, for example) can be managed and maintained within the secure confines of the corporate data center, and users still have their own

desktop environments. IT administrators can grant users controlled access to confidential information and intellectual property while protecting sensitive information.

VDI ADVANTAGES

With VDI, there is a unique environment for each and every user. Each of these environments can be completely customized with different applications and settings without impacting other users. Thus, users can be granted more control of their own virtual desktops, allowing them to install and modify applications if needed.

Moreover, VDI makes it possible to consolidate enterprise desktops. IT administrators can replace computer workstations with expendable, low-cost thin client terminals; this helps extend computer life cycles, as older computers can be repurposed for other duties.

With VDI, desktops are managed centrally, simplifying desktop installations, backups, and maintenance and reducing technical support and administration. And because user desktops—including operating system images—are running on servers within the walls of a centralized data center, data can be made more secure.

VDI decouples user desktops from specific hardware resources, enabling live migration of VMs between physical servers. Alternately, a VM can be suspended — the server can move the memory contents of the desktop VM to a disk; a VM can be provisioned on another physical piece of hardware (that is, a server); and the VM can then be brought back online. Once this

operation is completed, users can pick up right where they left off. It is also possible to have a session time out after a certain amount of time. The user's session can be disconnected from the client (still running on the server); the system can then suspend the session by moving the memory contents to disk and freeing up the hardware for another user. When the first user logs back in, the session is re-connected and the user can continue work—regardless of how long it has been.

VDI DISADVANTAGES

The flexibility and management benefits of VDI do come at a price. The appropriate infrastructure to support a VDI framework is comprised of enterprise-class server hardware connected to a storage area network (SAN). Even though server hardware is typically very robust and fault-resistant, it may become a single point of failure—the failure of the server may take down some or all client environments running on it.

This single point of failure extends to the network. VDI requires a low-latency connection between the client and virtual infrastructure; the client environments cannot operate without a network connection.

VDI PROVIDERS

VMware is the main provider of VDI software, though there are others. To build a VDI solution with VMware, IT administrators create a VM on the VMware ESX Server. IT administrators must then install a third-party VDI Connection Broker (such as ChipPC Virtual Desktop Center, Citrix Desktop Broker for Presentation Server, Dunes Virtual Desktop Orchestrator and Virtual Service Orchestrator, or LeoStream Virtual Desktop Connection Broker), then provision the VMware Networks Virtual Access Suite (VAS), install a desktop operating system on the VM (such as Windows® XP or Windows Vista®), and install desktop applications on the VM. Other VDI providers include Citrix, Surgient, and Virtual Iron (using software from Provision Networks).

AMD VALUE PROPOSITION

The performance of a VDI solution can be enhanced with AMD hardware. Large numbers of VMs mean large memory requirements. The AMD Opteron processor, with its increased memory addressability and high memory bandwidth, can support a large number of concurrent, memory-intensive VMs.

More VMs per Server

AMD Opteron processors with AMD Virtualization (AMD-V) technology are designed to make it possible

to run more VMs per server and to handle more users and transactions per VM than comparable server processors. AMD-V technology can help reduce overhead by selectively intercepting instructions destined for guest environments; it streamlines virtualization deployment, improves virtualization support, and helps guest x86 operating systems run unmodified at industry-leading execution speeds.

The *Direct Connect Architecture* helps eliminate the bottlenecks because everything is directly connected to the CPU—the processors, integrated memory controller, and I/O connect directly to the CPU and communicate at CPU speeds.

The *Integrated Memory Controller* is designed to improve performance on memory-intensive virtualization environments through high bandwidth, low latency, and scalable access to memory. The integrated on-die DDR2 DRAM memory controller offers available memory bandwidth up to 10.7 GB/s (with DDR2-667, which supports 667 million data transfers per second) per processor.

HyperTransport™ technology optimizes the movement of data and the sharing of resources among VMs for greater system scalability; it provides a scalable bandwidth interconnect between processors, I/O subsystems, and other chipsets, with up to three coherent HyperTransport technology links providing up to 24.0 GB/s peak bandwidth per processor.

Better Performance in Virtualized Environments

Rapid Virtualization Indexing, an enhancement to AMD-V technology, is designed to increase the performance of the most demanding virtualized applications through hardware-assisted memory management, enabling faster execution of guest code in many cases and lower hypervisor overhead in nearly all cases. Utilizing on-die silicon resources rather than software, Rapid Virtualization Indexing can greatly reduce hypervisor cycles and the associated performance penalty that is commonly associated with virtualization. Rapid Virtualization Indexing is also designed to decrease the “world-switch time”—the time spent switching from one VM to another—by 25% for greater application responsiveness.

Address space IDs (ASIDs) decrease context switching overheads and increase the performance of VMs.

Increased Responsiveness in Virtualized Environments

Unique to AMD Opteron processors, the *Tagged Translation Look-aside Buffer (TLB)* addresses switching times between VMs by mapping to the VMs' individual memory spaces. Competing solutions cannot distinguish one VM's memory space from another's, resulting in additional memory management overhead and reduced responsiveness when switching between VMs.

More Efficient Security

The *Device Exclusion Vector (DEV)* controls access to VM memory based on permission, isolating VMs for secure operation. The DEV performs these security checks in hardware, rather than in software, for greater efficiency; it creates Protection Domains that deny memory access for unauthorized requests from external devices, such as hard disks and network controllers.

Planning Ahead

A key factor in data center optimization is planning ahead for intelligent growth that enables faster deployment and better management of services. The AMD common-core strategy, same-socket infrastructure, and stable, long-term road map makes it possible to perform seamless upgrades of processor technology, helping IT administrators minimize the cost of transitions and maximize past investments in hardware, software, and personnel.

TERMINAL SERVICES

Terminal Services lets users access applications and data on remote, heterogeneous desktop environments over a local area network (LAN), wide area network (WAN), or other remote connections. Application execution takes place on the server, and the graphical user interface (GUI) is delivered to the client computing device via RDP. Terminal Services provides a cost-effective way to deploy line-of-business (LOB) applications that are frequently updated, hard to install, or need to be accessed over low-bandwidth

connections; it can also be used as a means to deliver the newest applications to an older PC population.

Terminal Services is somewhat similar to VDI, but instead of managing potentially hundreds—or even thousands—of centralized, virtualized, single user operating systems, IT administrators manage numerous user sessions in a single instance of the server operating system.

WHAT IS TERMINAL SERVICES?

Terminal Services provides remote access to a desktop through thin client software, allowing the client computer to serve as a terminal emulator. Terminal Services transmits only the GUI of the user session to the client. The client returns keyboard and mouse clicks (KVM information), which are processed by the server.

Users connect to the Terminal Server (the server component of Terminal Services) and see only their individual session, which is managed transparently by the server operating system. This individual session is independent of other client sessions. Client software can run on a number of client hardware devices, including Windows-based terminals, Macintosh computers, or UNIX-based workstations. Even some handheld devices running Windows Mobile® 6 are able to run a Terminal Services session.

With Terminal Services, the client appears as a window within the local desktop environment, and the client contains only the minimum amount of software necessary to establish a connection to the server and present the GUI. All other operating system functions, in addition to the applications, run on the server. Any input to the client is redirected over the network to the server, where all application processing takes place (see Figure 3).

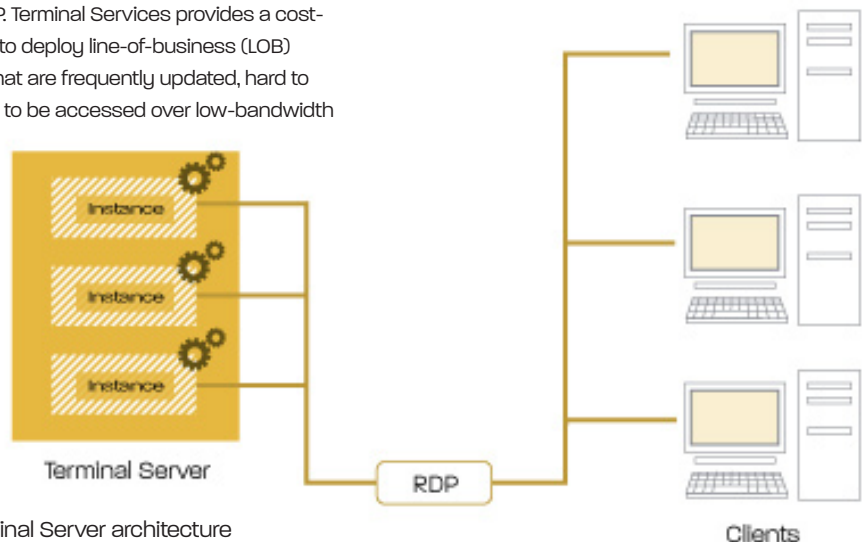


Figure 3 Terminal Server architecture

TERMINAL SERVICES

- Each user has a session on a single operating system instance that runs on the main server (one operating system to manage, many sessions)
- Applications execute on the main server
- Thin clients share processing resources

TERMINAL SERVICES USE CASES

Terminal Services provides a centralized application platform that can improve the desktop delivery in a number of scenarios. For example, in an organization without an assigned desk policy, users can work from different computers, using any available desk. If the computer on which a user is working does not have the necessary applications locally installed, the programs can be installed on the Terminal Server and made available; the worker can then use the programs just as if they were installed locally.

In the case of a corporate merger, companies often need to use consistent LOB applications on a variety of operating system versions and configurations. Alternately, companies outsourcing to partner organizations may need access to specific LOB applications, but not to the entire corporate network. Terminal Services provides a solution for both of these situations. Rather than going through the effort and cost of deploying all of the LOB applications to all of the computers in the merged company or the outsourcer, the LOB applications can be installed on a Terminal Server and made available as needed. This is especially useful if an application is difficult to patch or manage.

By securing an application and its data in a central location, it is possible to reduce the risk of accidental data loss caused, for example, by the loss of a laptop or mobile device. Centralizing applications and data ensures that as little data as possible leaves the corporate network. With Terminal Services, users, partners, or customers do not need to be granted full access to a company network or computers; they can be limited to even a single application, if needed, providing a means to control the IT environment.

TERMINAL SERVICES ADVANTAGES

The greatest advantage of Terminal Services is that users can run many desktop sessions on a single Terminal Server, and that server only has to manage one instance of the operating system. (Contrast this with VDI, where users have multiple copies of the operating system that need to be configured, managed, cleaned, and updated.) With Terminal Server, all program execution, data processing, and data storage occurs on the Terminal Server, centralizing the deployment of programs. Terminal Services can ensure that all clients can access the same version of a program. Because the software is installed on the server only once (rather than on every desktop throughout the organization), the

costs associated with updating individual computers are greatly reduced.

Terminal Services extends the model of distributed computing by letting computers operate as both thin clients and as full-featured PCs at the same time. Computers can continue to be used as they have been within existing networks, while also functioning as thin clients that are capable of emulating a Terminal Services desktop.

For example, if a company upgrades its SAP business solutions software, but not all of its computers can support the new SAP version, the company's IT administrators can create a Terminal Services session to support the computers that do not meet the requirements of the new version of SAP software.

Using a Terminal Server to run an application over bandwidth-constrained connections (such as dial-up or shared WAN links) is very effective for remotely accessing and manipulating large amounts of data because only a screen view of the data is transmitted, rather than the data itself.

TERMINAL SERVICES DISADVANTAGES

Though this can be an advantage because it can simplify management, one disadvantage of Terminal Services is that every client session on Terminal Server must use the same operating system. Using a single operating system diminishes the flexibility that IT departments can offer users who require unique or previous environments. Additionally, the Terminal Server must run several instances of an application within the single operating system, and some applications will not work in this fashion. The requirement to run several instances of an application can also cause complications for organizations using proprietary or customized applications.

Client operating system savings can be reduced by Terminal Server session licensing fees. Licensing in a Terminal Services environment can be complicated. If IT administrators opt for a Citrix solution, they will require a Terminal Services license in addition to a Citrix license.

TERMINAL SERVICES PROVIDERS

Terminal Services is a component of the Windows operating system (both server and client versions). Microsoft has a longstanding agreement with Citrix to facilitate sharing of technologies and patent licensing

OPERATING SYSTEM STREAMING USE CASES

Operating system streaming provides great benefits to call centers. The on-demand delivery of applications and the operating system to a diskless desktop client from networked storage makes it possible for call centers to give agents a desktop that can fully apply the sophistication of VoIP, while making centralized IT management a reality. With OS streaming, all programs execute locally on the desktop PC; this is critical to the success of the call center and ultimately the productivity of the agent, as the customer management software used in call centers often demands computing horsepower that only a full desktop can provide.

Moving to diskless workstations, which operate without the need for local persistent storage devices such as hard drives, RAM disk, or bootable CDs, can also be advantageous in classified work environments. Eliminating local storage at the desktop improves data security without complicating the day-to-day work of employees; networked storage provides the added benefit of data storage scalability.

OS streaming also provides benefits in education environments, where maintenance costs can be prohibitive. Most network maintenance costs—and most problems—stem from the need to configure and support every PC on an individual basis. Maintenance must be performed at the local PC hard drive, where the software and operating systems are installed. The situation becomes even worse in schools, where non-technical staff members are frequently charged with maintaining a complex web of hardware and software, and where students may unintentionally create a host of maintenance issues through the introduction of viruses, corrupted files, and spyware. Thin clients, which have no local processors or disk drives and rely entirely on the Ethernet cabling and server infrastructure for performance, may seem to be a good solution, but they have proven to be unable to power much of the rich, interactive multimedia applications frequently used in schools. Since OS streaming eliminates the disk drives but keeps the robust local processor, the network-bootable clients perform just like their traditional desktop counterparts, only without the local storage. A fresh configuration is provided at every startup. Nothing gets stored locally, so problems can be resolved at the server, eliminating the source of most desktop maintenance problems.

OS streaming can also make the introduction of new operating systems simpler. The transition to Windows Vista that many organizations will face over the next

few years can be made immeasurably easier with OS streaming. Organizations can maintain separate sets of Windows XP and Windows Vista disk images, for example. When users turn on their PCs, they can be presented with a simple menu that lets them easily choose which operating system they want to use for that session. They can switch between operating systems by simply rebooting and making a different choice. With OS streaming, IT departments can create tremendously flexible environments where they maintain a set of disk images that can be shared across their entire organization.

OPERATING SYSTEM STREAMING ADVANTAGES

OS streaming delivers performance. Since the client in the OS streaming scenario uses local resources to facilitate a user's session, there is very little difference between a locally installed copy of the operating system and the streamed version.

All clients boot from a single image. Therefore, as with Terminal Services, there is less patching and updating needed than in a traditional server environment. Moreover, a single fix can be propagated throughout an organization without a reboot. And because applications and data are stored and managed centrally, security and business continuity risks are mitigated.

OPERATING SYSTEM STREAMING DISADVANTAGES

In a WAN environment, IT administrators would need to set up an OS streaming segment on each major segment because of bandwidth requirements: a high speed LAN of at least 100 MB is needed to achieve strong client performance with OS streaming. The dependence upon a network connection for every client also means laptops cannot work offline.

To use OS streaming, IT administrators must rely on imaging technologies, which are prone to issues such as incorrect user names, "ghost connections," and missing drivers, which may impact performance.

OPERATING SYSTEM STREAMING PROVIDERS

Ardence, a Citrix company, has developed an operating system streaming solution that delivers Windows, SUSE Linux, Red Hat Linux, and Turbolinux operating systems, along with all their applications, on-demand from networked storage. Using the Ardence OS streaming technology, any x86-based computer—desktop, server, or device—can be provisioned, or re-provisioned, from bare metal to production in minutes. Ardence creates a vDisk (virtual disk, which can be private or shared) of

a desktop's operating system, applications, and configurations. This is done by taking a snapshot of the operating system and application image and storing the image as a file on the network so that the PC can boot from it. The computer starts running the assigned configuration immediately—there is no waiting to download the entire operating system.

Dell also provides On-Demand Desktop Streaming, an OS streaming solution that uses the Citrix Provisioning Server software to stream the operating system, applications, and data to desktop clients over a Gigabit Ethernet network. Users can take full advantage of the client's CPU and graphics processor, which provide a full desktop experience with application access and complete multimedia capabilities.

BLADE PC

Blade PCs—server-style blades that reside in a secure data center and hold the same basic components as traditional desktop computers—are gaining a foothold as a cost-effective way to conserve space and comply with security requirements in environments that may require multiple operating systems. Blade PCs combine the functions of a traditional PC with many of the benefits of blade servers. Unlike thin client technologies, Blade PCs support one user per discrete device; unlike Terminal Services, in which processing resources are shared, Blade PCs enable each user to connect to a dedicated processor, providing a more robust processing environment.

WHAT IS A BLADE PC?

With Blade PCs, the CPU, RAM, and a hard drive are in a “blade” form that plugs into an enclosure; they are entirely contained in a thin, modular circuit card placed in a centralized, secure location, such as a server rack. Many vendors have moved power supplies, cooling fans, and some management capabilities to the enclosure.

A Blade PC, in contrast to a thin client, is a complete computer that includes all the components normally found in a desktop PC, including the microprocessor, memory chips, hard drive, video card, and network card. Each server-side Blade PC contains the basic PC configuration: an operating system and the standard set of applications, patches, and drivers that the enterprise requires. The Blade PCs process the data and perform other server operations such as load balancing. The data itself is accessed from the data center, where it is stored in a user's individual account—a 2-GB partition, for example.

Blade PCs support one or more common operating systems, and they are intended to support one user per device. The user's monitor, keyboard, mouse, and other peripherals connect to a user connection device (or user port) that sits on the user's desktop and, in turn, connects to the Blade PC chassis.

Access to the device is usually achieved via RDP (or another protocol), which lets users log on to the Blade PC via a client device (usually a thin client). Once logged on, the user experience is largely the same as if the user was logged on to a local PC. A spare Blade PC or two can be kept in the enclosure for failover protection. Should a Blade PC crash while a user is logged on, the server will dynamically allocate another card to the affected user: logging back in takes just a minute or two. Some Blade PCs use an external physical layer fiber switch, which can cross-connect the blades to the various remote user connection desktop devices. This allows the IT administrator to map any port to any other port and dynamically allocate resources or immediately connect to a spare blade in the event of a failure.

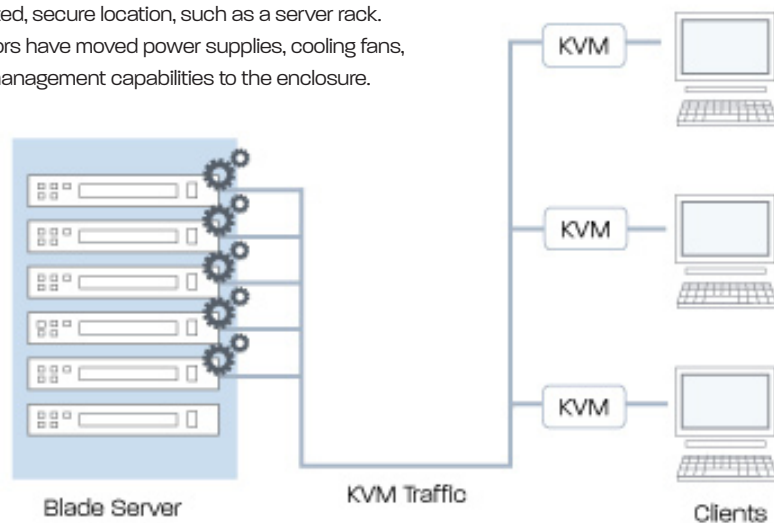


Figure 5 Blade PC

BLADE PC USE CASES

Blade PCs are ideal for call centers and billing/accounting departments, where users do not require mobile computers or powerful workstations, but can rely on an energy-efficient thin client. It may also be possible to reduce the number of desktops: an enterprise may employ 5,000 desktop users, but find that only 3,000 users log in during peak hours on average—thus there is no need to invest in 5,000 desktops.

Blade PCs are a preferred solution in secure, highly regulated environments, such as financial organizations, government installations, and other companies where security and manageability are important. A blade solution has no disk, processor, or Network Interface Card (NIC) at the desktop, so users cannot tamper with sensitive information. There have been numerous incidents where PCs or laptops containing customer financial data or mailing lists have been stolen. With Blade PCs, however, the blades reside in secure, not publicly available data centers.

BLADE PC ADVANTAGES

The TCO associated with Blade PCs has shown that it can be significantly less than that of traditional PCs residing at the desktop. The power requirement for the thin client can be far lower than that of an entire PC, so there is less wasted heat. Blade PCs also require less maintenance. Thin clients require less servicing, as they have no individual configuration; replacing one is as simple as plugging in a new unit. Replacement Blade PC cards, on the other hand, do need to be configured, but centralized maintenance makes it possible to image a new card and replace it in the rack, in less than an hour generally. The management software then adds the card to the pool. In addition to these savings, the ability to provide redundant blades leads to an increase in user productivity.

Blade PCs also offer a greater level of physical and virtual security compared to traditional PCs scattered throughout an organization; sensitive modules can be kept in the data center completely removed from the office work environment. And data residing in a central, secure location means a greater level of compliance with privacy standards. User access is controlled by passwords and central policy, and smart card or other identity systems can be integrated if necessary. Therefore, a company's greatest asset—its data—can be secure.

The processing/storage/network elements are separated from the user-oriented display elements; all of the heat, noise-producing, and power-consuming parts of the computer are kept in the data center in high-density racks far from the desktops, which can hold the quiet, space-efficient monitor, keyboard, and mouse.

An additional benefit of the chassis-based architecture is the relatively simple task of creating a dual-boot environment on each Blade PC that can provide, for example, a typical desktop environment during business hours and a Beowulf-clustered Linux environment (which could be integrated into an organization's grid computing fabric) during off hours.

BLADE PC DISADVANTAGES

While the TCO associated with Blade PCs is lower than even well-managed PCs, the acquisition and deployment costs are higher because of the associated infrastructure that is required.

In addition, RDP is less effective at delivering multimedia with Blade PCs, primarily because audio and video are not perfectly synchronized. In circumstances where there is increasing latency, there is a proportional decrease in the quality of the user experience.

BLADE PC PROVIDERS

Both Hewlett-Packard (HP) and IBM offer Blade PC solutions. HP provides a Blade PC that features AMD64 technology. The HP bc1500 blade PC is part of the HP Consolidated Client Infrastructure (CCI) product line aimed at reducing the cost of managing desktop PCs. The HP bc1500 blade PC features a low-power AMD Athlon™ 64 processor. The performance-per-watt of the AMD processors makes it possible for HP to install as many as 280 blades in a standard 19-inch rack. The CCI architecture, which uses RDP to communicate over the corporate network, also works with notebook and handheld PCs as clients.

ClearCube provides a PC Blade virtualization solution that combines VMware-enabled blades with management software, empowering organizations to virtualize the blades for comprehensive VDI solution. An additional vendor is Cubix, which provides the LaserBlade and LaserCubes.

CLIENT TECHNOLOGY CONSIDERATIONS

With so many options available for delivering applications to users, it can be difficult to determine which solution is best for a particular environment. The four technologies presented in this white paper are just a starting point—they can also be combined with each other and with other technologies for more possibilities.

In some situations, standard Terminal Services or a Citrix solution works best. In others, IT administrators could use VDI, the hosted desktops plus virtualization solution. In still other scenarios, a streaming technology such as Ardenance may be ideal, or IT administrators might implement a Blade-PC solution, such as the CCI solution provided by HP.

To determine which solution is best for their environment, IT administrators need to consider many factors, such as the application execution location, the need for a multi-user or single-user execution environment, and the installation of the applications. Other considerations include performance requirements, cost, and maintenance and administration capabilities. The following questions can help with the decision; these open-ended questions highlight the factors IT administrators should consider when determining how to best deliver desktops to users.

- Where does the application need to execute: locally on the client device or remotely in the data center?
- Do users need access from non-Windows clients?
- Are there high back-end data requirements?
- Does the application need to access data that is stored locally on the client device?
- Does the application need to integrate with other applications?
- Does the nature of the application lend itself to being made remotable?
- Do IT administrators need to ensure that data stays off the client device?
- What are the typical workloads run by clients/users?
- What is the network connection latency?
- What are the security/compliance requirements?

If IT administrators decide that the application will be executed remotely, then they need to decide what back-end architecture they will use.

- Can the application run on a multi-user system like Terminal Server?

- Will it need to be isolated onto a single user instance of Windows, such as in a Blade PC or VM?
- Does the user need administrative rights?

Once IT administrators determine where the application will run, they need to decide whether they will virtualize or isolate the application.

- Is it even technically possible to virtualize, isolate, or stream the application?
- Is there a need for dynamic application installs?
- Do all users need the application, or just some?
- How hard is the application to manage? Is it worth virtualizing?
- Is there a need for the raw application performance?

These questions help start the process of determining a solution. The best solution may be complex, however, and may require combining the architectures. For example, the IT administrator may want unique VMs for specific users, while still letting users connect to a pooled VM for a unique application that conflicts with others. Or an offshore developer may need a simple yet secure way of accessing the desktop that will not require the hosting company to install a great deal of software on the remote client. Combining VDI and Terminal Services may enable the IT administrator to address these issues and develop a unique solution that fits the organizations needs.

AMD TECHNOLOGY ADVANTAGES

At the core of all enterprise computing, regardless of the architecture used to deliver desktops, is the server and the server hardware. AMD is a leading manufacturer of processors, with a range of technology options that can be tailored to enhance the four architectures described in this paper, among others.

Power, Performance, and Efficiency

In many cases, AMD Opteron processor technology makes it possible for Terminal Services customers to support more users per server than other processors and increase the performance of demanding applications running on the same server, while helping reduce overall server and/or server farm power consumption.

Direct Connect Architecture, with its increased memory addressability and high memory bandwidth, can support a large number of concurrent users without affecting performance. Direct Connect Architecture also

provides improved PPW, which helps to reduce costs. This makes the AMD Opteron processor with AMD64 technology an ideal platform for OS streaming and Terminal Services solutions.

Add the “future-proofing” benefits of the AMD64 technology, which natively supports 32-bit applications while providing a smooth transition path to 64-bit computing, and current and prospective Terminal Services customers have a compelling value proposition.

The AMD Athlon™ 64 processor 1500+ is an ideal choice for Blade PCs. This advanced processor runs at 1 gigahertz (GHz) and has a mere 9-watt power envelope. This low-power envelope makes it possible to place as many as 280 blades in one industry-standard 42U rack, which can be more than twice the density of any other blade PC solution—key to consolidation in the data center.

The AMD Athlon 64 processor 1500+ is equipped with an integrated DDR (double data-rate) memory controller,

which lowers chip count, power consumption, and waste heat on each blade. The integrated memory controller also outperforms comparable off-chip memory controllers because the AMD Athlon 64 processor’s computational core communicates with the memory controller at its internal clock rate (which is faster than the external bus). It also has 128K of L1 cache and 512K of L2 cache to further speed overall performance. The power and performance of AMD’s latest technology makes Blade PC technology all the more inviting.

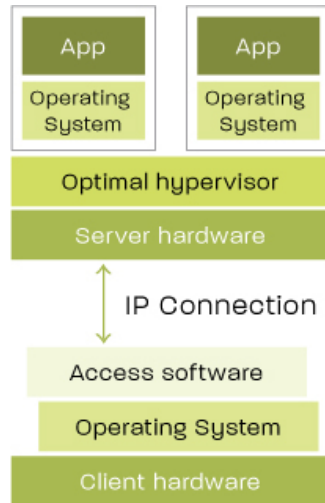
Enhanced Virus Protection

The AMD Athlon 64 processor 1500+ has Enhanced Virus Protection (EVP), which works in conjunction with the current releases of most operating systems to provide extra security against viruses and other malicious code that use buffer overruns. EVP is enabled by the operating system, and is now supported by current versions of Windows, Linux, Solaris, and BSD UNIX. With EVP, clients experience safer e-mail and Internet browsing, in addition to increased network and memory protection.

SUMMARY

Traditional enterprise PCs have served the needs of IT organizations well for many years. However, given the changing security environment, increased performance demands, and the need for greater flexibility at lower costs, it makes sense to rethink the traditional desktop and the enterprise client strategy. The alternative architectures currently available empower IT administrators to create the most secure, cost effective, and flexible environment possible.

In this white paper, four of these architectures were discussed: VDI, Terminal Server, OS streaming, and Blade PCs. Figure 6 and Table 1 show a comparison of these technologies. Each has its strengths and challenges, and each can provide performance and cost benefits in certain scenarios. While many, many more architecture alternative architectures exist for IT administrators to choose from, the focus in this paper is on these four more prevalent and relevant options.

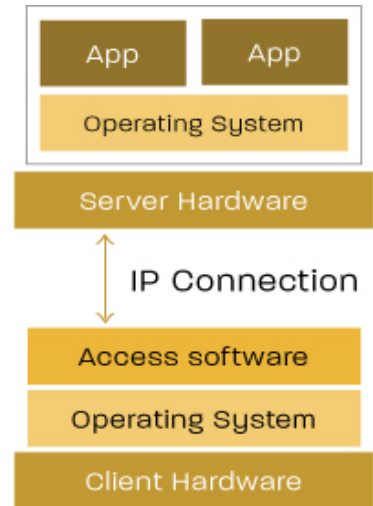


VIRTUAL DESKTOP INFRASTRUCTURE

Unique operating system

Shared hardware

- Desktop runs on virtualized server
- UI is local on thin client

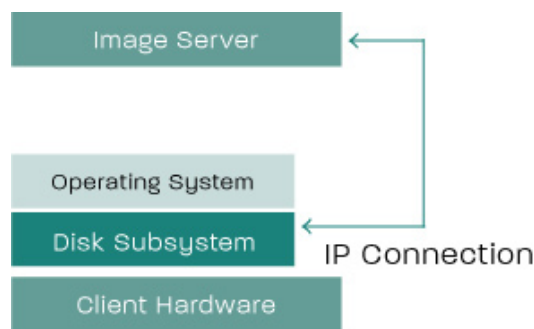


TERMINAL SERVICES

Shared operating system

Shared hardware

- Applications run on server
- UI delivered to client over network

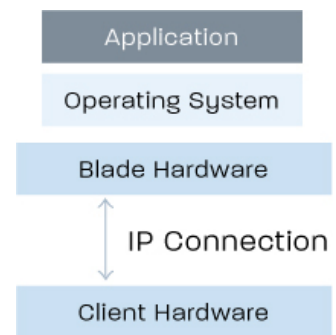


BLADE PC

Unique operating system

Unique hardware

- Desktop runs on the blade in data center
- UI local on thin client



OS STREAMING

Shared operating system

Shared hardware

- Uses SATA over IP
- Runs locally

Figure 6 Comparison of alternative architectures

Architecture	Operating system	Hardware	Where do the applications execute?	What travels between data center and client (or user)?	Summary
VDI	Unique	Shared (VMs hosted on main server)	Each user's VM	Access to VM	Many VMs and operating systems to manage
Terminal Services	Shared (each user has session on single operating system instance)	Shared	Central server	Screen refreshes	Thin clients share processing resources One operating system to manage, many sessions
OS Streaming	Unique	Blended (central storage, local PC)	Each user's PC (can be diskless)	Executable files (just enough binaries to execute applications)	Operating system on main, streaming server Relies on imaging technology, golden image is streamed
Blade PC	Unique	Unique (stored centrally)	Each user's centrally located Blade PC	Keyboard, video, mouse (KVM) information	Dedicated processor for each user User connects to Blade PC with RDP/thin client Supports one or more common operating system

Table 1 Alternative architectures

GLOSSARY

AMD64

A 64-bit superset of the x86 instruction set architecture (x86-64), the most commercially successful CPU architecture.

Back-end

A term used to characterize program interfaces and services relative to the user (either person or a program). A back-end application serves indirectly: supporting the front-end services either directly or through an intermediate program that mediates front-end and back-end activities.

Blade enclosure

Provides services such as power, cooling, networking, various interconnects, and management.

Blade server

Self-contained computer servers, designed for high density. Blade servers have many components removed for space, power, and other considerations, while still containing all functional computing components.

Bus

A subsystem that transfers data between computer components inside or between computers.

Chassis

The physical frame or structure of computer.

Connection broker

Controls session access to the server-side virtual desktops or other shared resources.

Data center

A centralized repository for the storage, management, and dissemination of data and information pertaining to a particular business.

Direct Connect Architecture

A computing platform available only with AMD64 processors (including the AMD Opteron and AMD Athlon 64 processors). With Direct Connect Architecture, there are no front-side buses. Instead, the processors, memory controller, and I/O are directly connected to the CPU and communicate at CPU speed.

Double-data-rate two (DDR2) memory

A random access memory technology used for high-speed storage of the working data. A type of

synchronous dynamic random access memory (SDRAM), one of many DRAM implementations. The primary benefit of DDR2 memory is the improved bus signaling and the operation of the memory cells at half the clock rate (one quarter of the data transfer rate). DDR2-xxx (as in DDR2-667) denotes data transfer rate. Bandwidth can then be calculated by taking transfers per second and multiplying by eight (because DDR2 memory modules transfer data on a bus that is 64 data bits wide, a byte comprises 8 bits—8 bytes of data per transfer).

Failover

A backup operational mode, used to make a system more fault tolerant, in which the functions of a system component (a processor, server, network, or database) are assumed by secondary components when the primary component becomes unavailable (through either failure or scheduled downtime).

Front-side bus

The physical bidirectional bus that carries all electronic signal information between the central processing unit (CPU) and the northbridge (also known as the memory controller hub, one of the two chips in the core logic chipset on a PC motherboard).

Front-end

A term used to characterize program interfaces and services relative to the user (either a person or a program). A front-end application is one that users interact with directly.

Graphical User Interface (GUI)

A graphical (rather than purely textual) user interface to a computer.

Gigabit Ethernet

A transmission technology based on the Ethernet frame format and protocol used in LANs. Gigabit Ethernet provides a data rate of 1 billion bits per second (one gigabit). Gigabit Ethernet is defined in the Institute of Electrical and Electronics Engineers (IEEE) 802.3 standard, and it currently serves as the backbone of many enterprise networks.

Independent Computing Architecture (ICA)

A thin-client protocol from Citrix Systems.

Internet Protocol (IP)

The method (protocol) by which data is sent from one computer to another on the Internet.

Kernel

The core of a computer operating system, providing basic services for all other parts of the operating system. The kernel can be contrasted with the shell, or outermost part of the operating system that interacts with user commands.

Latency

The time between the initiation of a request for a byte or word in memory and the time it is retrieved.

Load balancing

Used to spread work across two or more computers, network links, CPUs, hard drives, or other resources, in order to optimize resource utilization, throughput, or response time.

Local Area Network (LAN)

A group of computers (and associated devices) that share a common communications line or wireless link. Typically, the connected devices share the resources of a single processor or server within a small geographic area.

Network Interface Card (NIC)

A computer circuit board (card) that is installed in a computer so that it can be connected to a network.

Quad-core

Combines four independent cores onto a single die or integrated circuit (IC).

Remote Desktop Protocol (RDP)

A multi-channel protocol that allows a user to connect to a computer running Terminal Services. Clients exist for most versions of Windows and other operating systems such as Linux, FreeBSD, Solaris, and Mac OS X. The server listens, by default, on TCP port 3389.

Server farm

A group of computers (servers) working together and housed in a single location. The servers may have individual operating systems or a shared operating system, and they may be set up to provide load balancing (when there are many server requests) or fail over (if one server fails, another can act as backup).

Storage Area Network (SAN)

An architecture that is used to attach remote computer storage devices (such as disk arrays and tape libraries) to servers in such a way that the devices appear to be attached locally to the operating system.

Thin client

A client computer (or client software) that depends primarily on the central server for processing activities and instead focuses on conveying input and output between the user and the remote server. (A thick, or fat, client does as much of the processing as possible, only passing data to the central server for communications and storage.)

Virtual Machine (VM)

A complete image of a PC, including its applications, operating system, and device configurations, all encapsulated into a set of files that can run on any PC.

Wide Area Network (WAN)

A geographically dispersed telecommunications network (as opposed to a LAN).