ENTELECHY
ASSOCIATES

# VDI Appliances Accelerate and Simplify Virtual Desktop Deployment

## Introduction

Early desktop virtualization adopters quickly learned that success or failure of a VDI implementation very often hinged on the ability to deliver the right storage infrastructure; a fact that many industry commentators and vendors quickly picked up on. So much so that, in many peoples' eyes, the most important infrastructure design consideration has become the need to deliver IOPS – Input/Output operations Per Second. There is no doubt that storage is a critical part of any successful VDI system; however, this understanding does not paint the whole picture. The introduction of dedicated VDI appliances creates an opportunity to rethink VDI infrastructure design from the ground up, simplifying and accelerating delivery. At the same time, variation in VDI appliance design means that attempting to gauge appliance performance by reviewing processor and IOPS specification alone is insufficient to gauge likely system performance.

The VDI appliance creates an opportunity to rethink VDI infrastructure design from the ground up, simplifying and accelerating delivery

## The Importance of IOPS to VDI Success

Desktop operating systems like Microsoft Windows are designed to run with dedicated system resources – most importantly memory and storage. Ensuring adequate resources are available on a standalone desktop PC is never a challenge. A PC can readily accommodate 16 GB of memory and conventional low-cost (7,500 RPM) spinning disk storage can offer in excess of 1 TB and 80 IOPS per spindle, which is more than adequate for Windows 7.

In a VDI environment, where many desktops are concentrated onto a single hypervisor platform, meeting these requirements becomes more difficult. Storage capacity is never a problem, compared to many systems, the amount of disk storage required for a VDI environment is low. Attaining the required throughput is, however, a different matter. With many hundreds, and in some cases many thousands, of desktops all sharing a common storage infrastructure, demand for IOPS presents a significant challenge. Furthermore, as a desktop operating system, Windows 7 was developed with the assumption that all operating system resources are available for its exclusive use. With VDI, many Windows desktops simultaneously and independently read and

write to a shared set of storage resources in a manner that looks highly random since there is no coordination among the desktops. Random reads and writes present an interesting performance challenge for latency sensitive disk-based storage infrastructure.

The impact of this problem is most dramatically demonstrated in the “boot storm”. The Windows boot and logon processes generate many times more IOPS than steady-state user operations. When many desktops are all starting at once, for example at the start of a call center shift, the storage infrastructure will see an order of magnitude increase in load compared to the rest of the day. In poorly specified systems, this boot storm will overload the storage infrastructure, starving Windows desktops of resources, resulting in excessively long startup times and degrading performance across all active systems.

Conventional storage solutions attempt to meet the IOPS challenge by scaling out – deploying many more high-performance disk drives than are needed for capacity purposes to deliver the required IOPS. This approach is capital intensive, represents a significant ongoing operational cost and is ultimately unsustainable as virtual desktop numbers increase. Scale-up solutions make use of Solid State Drive (SSD) storage systems which achieve much greater IOPS to meet boot storm load levels. However, SSDs have relatively low capacity and are expensive in comparison to conventional drives. Server based high-performance flash memory controllers can be used in conjunction with disk-based SANs to increase available IOPS, but these often cannot be used with blade servers which are frequently cited as the preferred hardware platform for VDI environments.

## Adopting a holistic approach

If more IOPS delivers better performance, then isn't the secret to success to design storage infrastructure for maximum IOPS? On the surface the answer is simple, yes. However, if higher performance can only be achieved through increased cost and complexity, it is time to explore alternatives.

One of the primary reasons CTOs use to express the need to adopt VDI is to achieve full control of the desktop. The motivation behind this goal often varies, it may be to revoke administrator-level control to prevent unauthorized changes and to secure enterprise data, or replace the desktop PC

with a thin-client or tablet. It is slightly ironic then that by looking back to the desktop PC it is possible to gain a better view of how best to meet the IOPS challenge.

In many respects, the benchmark for balanced systems are conventional desktop PCs. With over 30 years of working together to guide their progress; processor, memory, and storage technologies have achieved balance. A PC combines all these components into a single integrated system. Processor, memory, and storage are all in the same box, together with the desktop operating system and applications. This close physical integration means that when a call is made by the operating system, the response from storage can be measured in microseconds and the system will deliver exactly what is expected of it.

Buying an off-the-shelf PC brings with it a high degree of confidence that all components will be balanced for both performance and budget. It's no coincidence that manufacturers today tend to specify 2 GB of memory as the base configuration for a PC with an Intel i3 processor, but increase that to 4 GB with an i7 processor in anticipation of the greater demands that might be placed on it. Of course some people still find the need to maximize performance by sourcing individual components and bringing them together in a custom build, but 'science project' PCs like this seldom offer good value and are always subject to the risk that the individual parts will not work well together.

Taking the PC as an example of best practice in balanced system design, the key to designing the storage infrastructure of a VDI environment is not as simple as maximizing IOPS; it is rather in achieving balance between the compute resources and storage. The goal is not to ensure that the storage infrastructure can deliver the highest possible level of throughput, but to ensure that it is balanced to accommodate the workload placed upon it. Bringing together multiple components from multiple vendors to achieve maximum performance is certainly possible, but it can place a significant burden on the IT department to integrate these components and subsequently support them. Just as science project PCs are best left to the hobbyist, so science project VDI environments have no place in the enterprise.

# VDI Performance Considerations

Overall performance of correctly-sized VDI environments is largely governed by the combined network latency between the endpoint and data center server, and between the server and the storage system. Endpoint network latency affects the responsiveness of the user desktop; as latency increases above 100-150 ms, mouse and keyboard response lag becomes increasingly apparent. Keystrokes do not appear on screen promptly and mouse cursor movement lags behind user input.

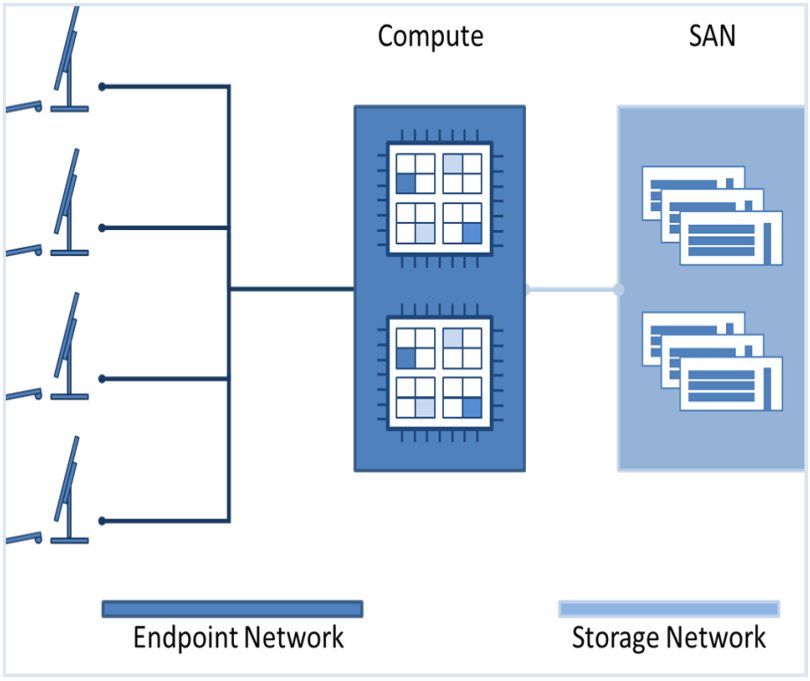


Figure 1. Basic VDI implementation using separate Compute and Storage components

Storage network latency can be a much bigger issue. As has been shown, desktop operating systems, such as Microsoft Windows, are designed to run on dedicated storage with high bandwidth and low latency. Adding even a few milliseconds latency between compute and storage can greatly reduce IOPS and visibly impact desktop performance, making the system appear slow and unresponsive. The obvious solution is to deliver the same low latency storage service that a conventional desktop PC can provide, but do so on a shared server in the data center. Achieving this with conventional SAN technology is possible, but only by introducing expensive caching technology or introducing additional third party components that add risk and complexity to the design.

## All VDI environments have the following five core elements:

- Endpoint: The device (PC, laptop, thin-client, tablet, etc.) from which a user accesses the virtual desktop
- Endpoint Network: The network that connects the endpoint to the virtual desktop environment
- Compute: The server hardware that hosts the virtual desktop
- Data Center Network: The network that connects the compute environment to the storage environment and other critical applications within the data center
- Storage: The system that stores end user data. Frequently a SAN with cache acceleration

# The VDI Appliance

An increasingly attractive alternative to the science project VDI environment, is to consider the VDI appliance – an integrated platform that combines compute, storage, and virtualization software as a single platform that eliminates much of the complexity of implementing VDI. By close coupling the compute and storage elements into a single appliance, the storage network becomes part of the appliance, eliminating any latency between compute and storage and providing a virtual desktop that matches or even exceeds the performance of a conventional desktop. At the same time, as an integrated solution, the VDI appliance hardware is chosen to ensure that it appropriately matched for its "rated" workload so any risk associated with selecting and sizing components is effectively eliminated.

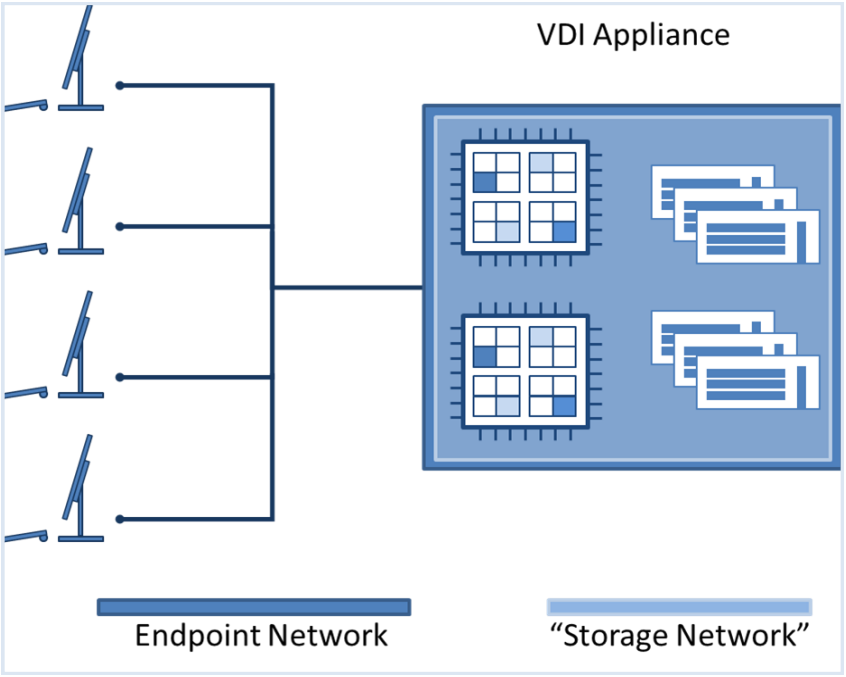

Figure 2. VDI appliance with integrated Compute and Storage components

VDI appliances are offered in multiple configurations optimized to support a given number of desktops, making it possible to scale up and scale out as necessary to support increases in capacity as demand for virtual desktop services grows. Solutions range from branch office platforms such as Cisco's Integrated Service Router (ISR) capable of supporting only 20-25 desktops, through to full-size VCE VBlock systems capable of supporting thousands of concurrent users. Further differentiation of VDI appliances extends to high availability and fail-over capabilities, and to the hypervisor and VDI platforms supported. All VDI appliances share a common concept of hosted virtual desktops sharing local compute and memory, but have dramatically different approaches to storage, both in terms of overall storage architecture and in delivering the IOPS needed to meet performance expectations. The diversity of approaches the vendors have adopted, and the wide variation in size of different systems, make direct comparison based on component specification difficult; however by considering each solution as a whole, valid comparisons are possible.

# Assessing VDI Appliance Performance

The introduction of VDI appliances offers a new way to consider VDI performance. Instead of focusing on the performance of the discrete building blocks that create the VDI appliance (i.e., processing power, memory capacity, and storage IOPS), it is better to measure what the system as a whole delivers.

Again, we can return to the PC for guidance. Considering just the breadth of choice available in processor types – overall architecture, number of cores, cache size, processor clock speed, front-side bus speed, etc. it is beyond most people's ability to assess which combination of factors will deliver the best performance for a particular use case. It is possible to measure individual component performance, but when taken in isolation these performance figures are of limited value. Instead, PC performance is evaluated using widely available benchmarks, making the results more accessible and understandable by all.

The same approach can be taken when it comes to assessing the performance of VDI appliances. A VDI appliance delivers desktops, and it is this ability to deliver desktops that should be assessed, not which processor it uses or how many IOPS its storage system delivers.

Two primary tests are needed to assess a system's performance:

- Start-up performance – measurement of the time needed to start the "rated" number of sessions
- Overall system performance – a conventional benchmarking test to validate the rated number of sessions is feasible

Tests should be carried out for both a single appliance, and for the maximum number of appliances that the vendor recommends grouping in a single service domain or "cluster" – this is to assess if scaling out to multiple appliances impacts performance in any way.

In addition to basic performance testing, consideration should be given to testing system performance in the event of appliance failure. Should an appliance fail, the default behavior of any VDI system would be to attempt to re-allocate sessions across the remaining appliances. Understanding the operational characteristics of an appliance when it is attempting to host greater than its recommended number of desktops will provide valuable data when designing for disaster recovery and business continuity conditions.

Assessing how the system performs when recovering from failure may also be needed if system performance guarantees will not permit any performance degradation during system recovery.

## Login VSI is widely accepted as providing a technology neutral means of assessing the performance of VDI systems.

There is already a de facto standard for assessing the performance of VDI systems – Login Consultants Virtual Session Indexer (Login VSI). Login VSI is widely accepted as providing a technology neutral means of assessing the performance of VDI and Remote Desktop Services (RDS) systems. Following this lead, adopting Login VSI as a standard testing solution for all VDI appliances is recommended over proprietary testing methodologies.

## Conclusion

VDI appliances offer a different perspective on implementing hosted virtual desktops. Instead of focusing on software and professional services to integrate the hardware; VDI appliances deliver an integrated package designed for rapid deployment that takes the guesswork out of specifying, costing and deploying VDI platforms.

These appliances are viable in any scenario from a small branch office without on-site support, to large enterprise data centers supporting thousands of users.

Individual solutions differ in cost, capacity, and features, but all provide the same core service of a scalable VDI platform designed for predictable performance and operation. List price per desktop varies significantly between solutions, but each is an order of magnitude lower than that of first generation VDI solutions.

## About Entelechy Associates

Effective desktop and application delivery architecture design and product selection is both difficult and complex. Entelechy Associates helps clients reduce risk while accelerating project implementation by providing in-depth, hands-on experience in designing enterprise-class desktop and application delivery environments. Our services are designed to support clients during all aspects of the requirements management, planning, and decision-making processes. Resulting in the best possible alignment between business need and strategic IT service delivery.

For more information on Entelechy Associates services see http://entelechy-associates.com/services

## Disclaimer

The terms and conditions of product sales are limited to those contained on the Entelechy Associates website at Entelechy-Associates.com. Notice of objection to and rejection of any additional or different terms in any form delivered by customer is hereby given. For all products, services and offers, Entelechy Associates reserves the right to make adjustments due to changing market conditions, product/service discontinuation, manufacturer price changes, errors in advertisements and other extenuating circumstances.

Entelechy Associates and the EA logo are trademarks of Entelechy Associates llc. All other trademarks and registered trademarks are the sole property of their respective owners.

This reference guide is designed to provide readers with information regarding the virtual workspace. Entelechy Associates makes no warranty as to the accuracy or completeness of the information contained in this reference guide nor specific application by readers in making decisions regarding client computing. Furthermore, Entelechy Associates assumes no liability for compensatory, consequential or other damages arising out of or related to the use of this publication. The content contained in this publication represents the views of the authors and not necessarily those of the publisher.

## License