Memory and Storage Best Practices for Desktop Virtualization

Balancing User Experience, Cost and Flexibility

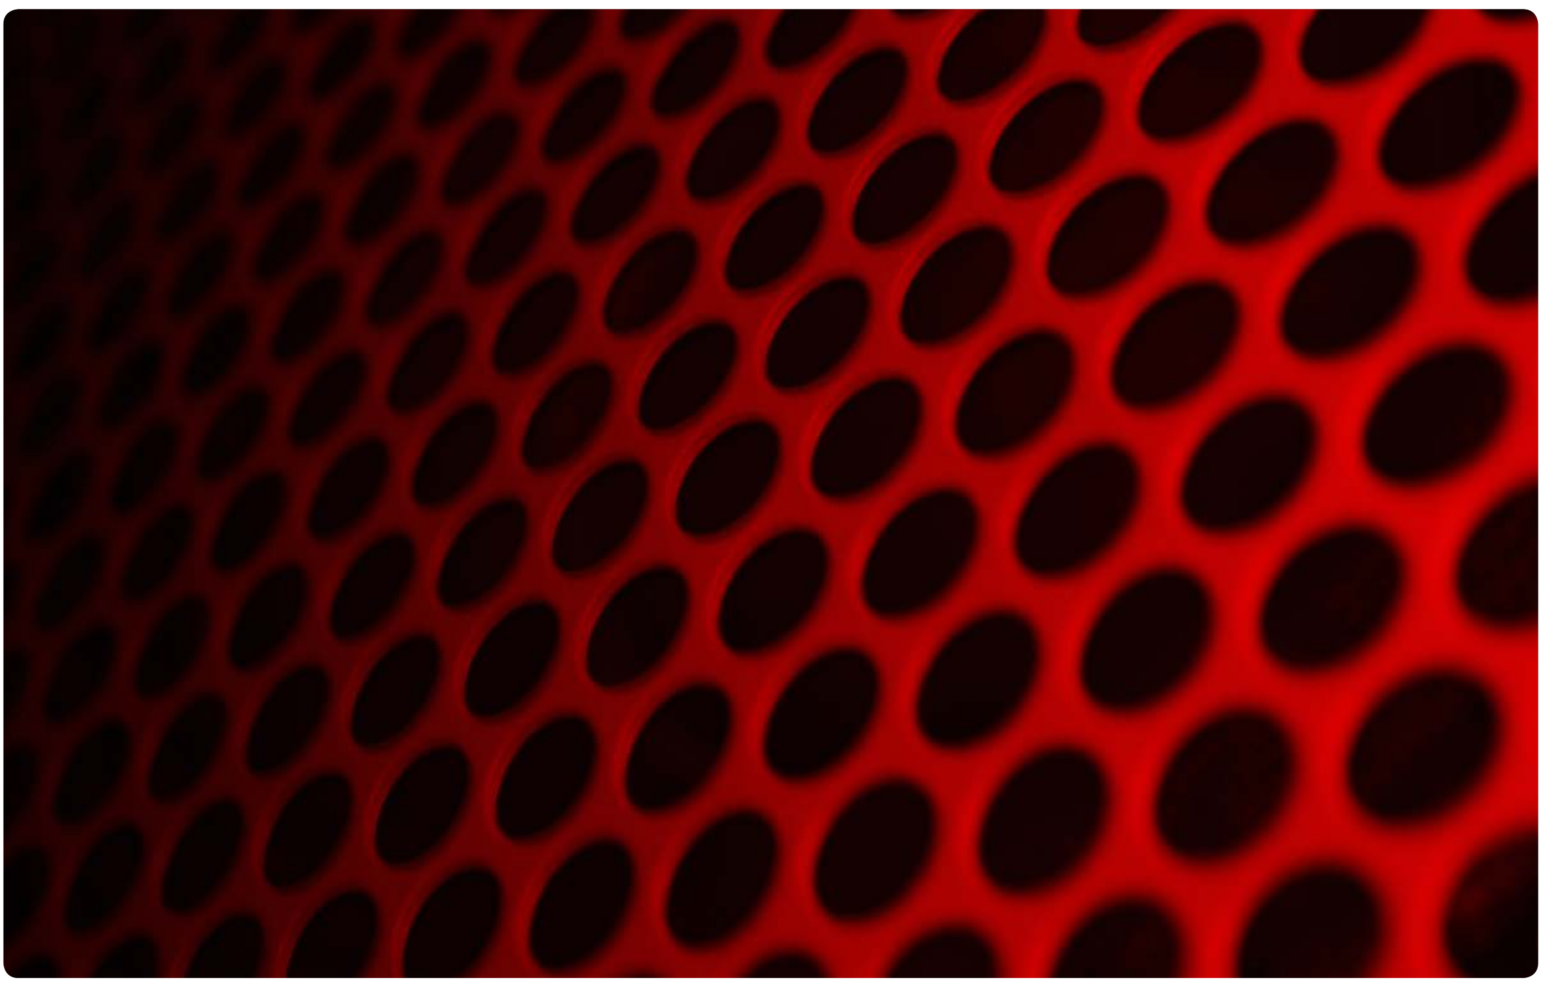

Kingston TECHNOLOGY

Table of Contents

Introduction and Summary

The current approach to managing desktop computing workloads has remained relatively unchanged since the introduction of the personal computer some thirty years ago. There's little doubt a significant desktop transformation is currently underway. This change is about creating a new, agile, available and optimized workspace. One that increases end user productivity while increasing management and operational attributes. This change involves a radical paradigm shift—one that both consolidates the compute workload and creates significant efficiencies in the way end user resources are provisioned, optimized and managed.

This paper will focus and discuss virtual desktop infrastructure (VDI); a desktop delivery approach that, when implemented properly, offers operational efficiencies not realized in current desktop computing support practices. Specifically, VDI is an evolving art that is rapidly becoming more science, or metrics driven.

Founded in 2009, Liquidware Labs™ is a provider of solutions for next-generation physical and virtual environments. Our technologies have been described by analysts as the industry's first "on-ramp to VDI," providing a complete software suite that enable organizations to realize the benefits mentioned above. Liquidware Labs has been involved in this desktop transformation for the last five years. Our industry experience, gathered in the trenches and with alliance partners such as Kingston Technology Corporation, enable us to share the best practices contained in this paper.

With a new approach, a new set of tools and metrics must be defined to ensure appropriate end results are realized. This paper will share some observations and best practices to provide you with the visibility required to balance end user experience, project cost and operational efficiencies. We will examine how overall resources like memory and storage contribute to a successful desktop transformation strategy. Additionally, we will discuss the critical need to quantify end user experience as a means to deliver the right resources at the right time.

Observations From Desktop Virtualization Projects

Since its inception, Liquidware Labs has been heavily engaged in the transformation of the desktop. With visibility based upon leveraging our Stratusphere FIT product to support assessments, and Stratusphere UX to validate performance, we have a long history that encompasses over 400 metrics-driven desktop virtualization projects. Agnostic of approach, the results of this history are unambiguous—success in VDI requires a new approach to how server-class resources are sized and implemented to support desktop workloads. Whether you invest upfront with an assessment or jump in and optimize later, there are critical attributes that must be considered.

Many early adopters of VDI assume that measuring desktop workloads prior to beginning a virtualization project serves only to support the sizing and build-out of the host server environment. While this is partially true, there are other important benefits associated with this step, specifically as it relates to creating the optimal virtual machine (VM) image. We believe there are a couple of primary benefits to a metrics-driven approach to VDI; namely:

- Capture the baseline user experience in the current environment—this is critical to ensure you provide an equal-to-better user experience when physical desktops are converted to virtual. Related, this step provides the ability to perform a before-and-after comparison of resource use and ultimately, end user experience.
- Monitor application use as it relates to desktop pools and images—if for no other reason, this attribute of the assessment helps you to better understand what applications are used versus installed. This benefit also provides visibility into user and group resource requirements.

Regardless of when you measure, do not miss the critically-important step to quantify end user computing resource requirements to support and ensure optimal end user experience. It is the cornerstone of a successful and optimized desktop virtualization implementation. This paper will not detail the specific steps or process with respect to assessing user and machine-specific workloads. Rather, this paper will highlight the importance of memory and storage resources; specifically the critical role each plays in the end user experience and overall performance of virtual desktop workloads.

Why Do Server Memory and Storage Resources Matter?

Delivering end user computing resources in a virtual desktop architecture is profoundly different from how we've provisioned, managed and optimized desktops in the past. More to the point, the way we manage server and storage resources to support VDI is about identifying and minimizing resource bottlenecks in your environment.

We have found that the most common resource bottlenecks observed while leveraging the Stratusphere UX product relate to consolidation ratios, memory and storage resources. These are very common performance occurrences, which can be prevented during the early phases of implementation.

- Poor consolidation ratio—this very common bottleneck is due to unbalanced resource usage in host servers. Understanding how CPU and memory play a role in optimizing VDI performance is critical to meeting total cost of ownership (TCO) and return on investment (ROI) goals.
- Improperly sized VM memory—memory-to-disk swapping on the guest OS is another common issue that can cause end user performance issues in a VDI architecture. This can be especially tricky as host memory page sharing and ballooning do not prevent swapping if the guest VM OS "thinks" it's near capacity.
- Storage and "boot storms"—successful VDI deployments also minimize the number of VM images required to satisfy all use cases. This desire can have the negative consequence of creating boot, or login storms; especially if storage requirements are not measured and allocated for both average and peak requirements.

Memory Sizing for the Virtual Machine Client and Host

Estimating memory requirements for virtual desktops is difficult. The architecture can be complex and ever-changing user habits contribute to the need to closely monitor virtualized resources. For example, if the random access memory (RAM) allocation is too low, storage input and output performance can be negatively affected as too much memory-to-disk paging occurs. Conversely if the RAM allocation is too high, storage capacity can be negatively affected as the paging file in the guest operating system, and associated swap and suspend files for each VM, can grow too large.

When preparing or reviewing how memory resources are assigned to the virtual environment, there are three main objectives to keep in mind:

1. Swapping is the enemy—memory-to-disk paging in a virtual environment is a performance killer. You must measure and provision memory in the virtual guest operating system to minimize paging.

2. Balanced host build—care must be taken to balance the VM memory with physical host memory to ensure you don't reach maximum memory capacity when CPU allocations are low.

3. The right resource for the right user—associating users with appropriate VM templates, which are adapted to their needs, will optimize resource usage and ultimately reduce the overall solutions return on investment (ROI) and total cost of ownership (TCO).

Understanding How Memory Is Delivered and Used

Whether designing a new VDI or adjusting an existing infrastructure, one critical factor to consider is how virtualization affects memory calculations. Everything from the hypervisor of choice to the virtual desktop approach can have an influence on memory usage. Specific factors to be considered and applied include:

- Hypervisor operating system overhead
- Hypervisor page sharing and memory ballooning
- Operating system version and CPU instruction set
- Application delivery (application virtualization can have an influence on guest VM memory)
- Desktop virtualization approach (user sessions vs. virtual desktop, persistent vs. non-persistent)

For example, Microsoft Windows XP will begin paging memory to disk when the majority of guest VM memory has been consumed. As noted, memory-to-disk paging in a virtual environment can have a significant impact on performance. Therefore, memory allocation should be adapted accordingly: increasing memory allocated to the VM, while reducing the swap space typically assigned to physical environments. You can't be too generous with respect to allocation of memory to each VM in an effort to avoid disk paging.

This approach doesn't necessarily come at the expense of consolidation and overall cost when you factor in hypervisor memory management features, such as over allocation and page sharing. Assuming that all desktops share the same operating system, you can safely predict that approximately 50 percent of memory allocated to clients will be shared and consumed on the host.

Capturing Baseline Memory Usage For Guest VMs

To determine overall memory requirements, we recommend using metrics captured at the user or machine level. If available, the use of average and peak average data will allow you to properly capture the total amount of memory consumed. While RAM peaks are common, we recommend normalizing this metric for sizing, as building for peak often results in oversizing. Stated another way, peak average results represent the "average" usage for a particular user/machine over a specific time period—this is typically a good "high water mark" for sizing purpose. See the figure above for observable differences between average, peak average and peak RAM consumed.

Metrics	Allocated	Average Used	Peak Average Used	Peak Used
CPU MHz	99.11 GHz	12.42 GHz	19.14 GHz	47.7 GHz
RAM	94.9 GB	49.99 GB	51.12 GB	66.57 GB
Page/Swap File	92.1 GB	10.96 GB	11.31 GB	23.83 GB
Storage Capacity	3.59 TB	2.02 TB	2.02 TB	2.3 TB

Grouping Users By Desktop Profile

Grouping users by their associated desktop profile is an important consideration as it has benefits relating to management and resource utilization. The primary goal in this exercise is to match each user (or user group) to its appropriate profile. In this context we define a desktop profile as the method of delivering the workspace to a user—it could be a VM, a terminal session or even a physical desktop.

There are many factors that play a role in determining the number of desktop profiles and how they are assigned to users; they can include user location, user role or function within the organization, travel or security policy, to name a few. User segmentation by profile and how those decisions contribute to overall end user management is beyond the scope of this paper. We'll be examining how those decisions effect memory resources, and ultimately server host memory.

This exercise begins with a thoughtful approach to defining profiles. As noted, this process can be challenging for those unfamiliar with virtualizing end user environments. Most common implementations include between two and five profiles—striking a good balance between user segmentation and manageability. The typical VDI environment might include the following profile groups:

- Default/basic virtual desktop
- Advanced virtual desktop
- Terminal Services/XenApp session, Mirage user, etc.
- Non candidate/physical desktop/laptop or other delivery

With the exception of implementations limited to one group of users (a call center, for example), it would be the norm for all end user virtualization efforts to have two desktop profile groups, at a minimum. This statement is made solely on statistical observations and how typical organizations' users naturally segment based on resource requirements.

For Those Not Yet Virtualized

Next we will demonstrate how a balanced approach to profile grouping can be determined. This exercise requires prior analysis in the form of an assessment—without these metrics in hand, quantitative user profile segmentation would be exceptionally difficult. The figure below is output from Stratusphere Preview Inspector, VDI FIT Consumption–Memory. You should immediately note that users' naturally fall within a few groupings based upon consumed memory.

Where these natural groupings segment is called an inflexion point; specifically, where the amount of memory used defines clear desktop profile groupings. In the example below, it is clear that the majority of users are consuming under 1000MB (1GB) of memory. A second smaller group can be realized between 2000MB (2GB) and 3000MB (3GB) of memory used. And finally a few user exceptions exceed 4000MB (4GB) of memory used. This type of trending and data forms the starting point for casting your profile groupings.

Employing a thoughtful approach to grouping users into appropriate desktop profiles can be tricky. There are many other variables to consider; however, gathering quantitative metrics will go a long way to optimizing your resources and creating a manageable environment. For those that may be looking to optimize available memory resources in an existing VDI environment, please note the guidelines described above will still apply, but the approach requires a different path.

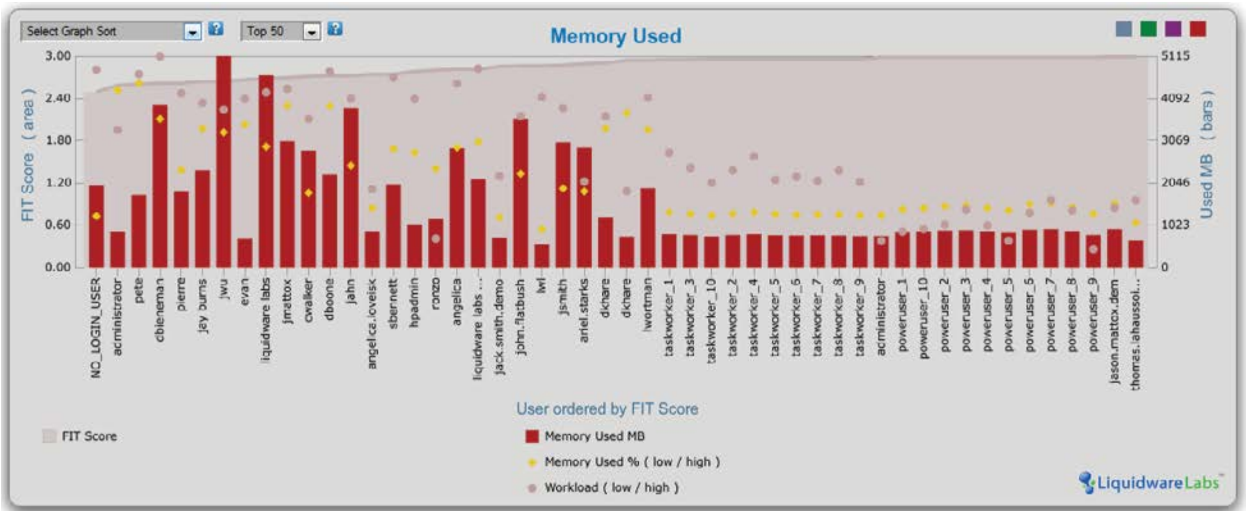

For Those Looking to Optimize an Existing VDI Environment

As noted, end user performance in a virtual environment can be a balancing act between memory and storage—with the specific goal of minimizing memory-to-disk paging. The metrics shared below are dashboards from Stratusphere UX. Please note that we recommend you first turn your attention to memory and page faults. In the below example, memory allocations are causing some machines to experience virtual memory-to-disk paging (page faults). You will recall, these are one of the primary contributors to a poor end user experience.

User	FIT Score	UX Score	MEMORY USED %	Memory Allocated MB	Memory Used MB	Page Allocated MB	Page Used MB	Pagefile Used %	Soft Page Faults	Hard Page Faults
pete	B+	B-	85.7 %	2,010 MB	1,724 MB	2,056 MB	958 MB	46.6 %	1,850	155
administrator	B+	B-	78.8 %	1,024 MB	807 MB	1,625 MB	1,108 MB	68.2 %	1,594	13
evan	B+	B-	78.0 %	1,015 MB	792 MB	2,042 MB	1,322 MB	64.8 %	2,431	
jmattox	B+	B+	76.2 %	3,980 MB	3,031 MB	4,818 MB	1,587 MB	32.9 %	1,702	260
dboone	A-	B+	73.3 %	2,934 MB	2,150 MB	2,933 MB	945 MB	32.2 %	3,094	129
dkhare	A+	B-	73.1 %	1,023 MB	748 MB	768 MB	420 MB	54.6 %	361	18
dbieneman	B+	B+	72.7 %	5,609 MB	4,076 MB	5,609 MB	1,180 MB	21.0 %	1,420	198
lwortman	A+	B+	65.8 %	2,934 MB	1,929 MB	2,933 MB	633 MB	21.6 %	1,206	79
dkhare	A+	B+	63.1 %	1,753 MB	1,106 MB	1,638 MB	375 MB	22.9 %	1,062	54
liquidware labs inc	A	B+	58.1 %	3,578 MB	2,078 MB	3,577 MB	86 MB	2.4 %	1,198	164
sbennett	B+	B+	58.0 %	3,578 MB	2,074 MB	3,577 MB	74 MB	2.1 %	1,339	93
liquidware labs	B+	A	57.1 %	8,126 MB	4,636 MB	8,126 MB	127 MB	1.6 %	3,902	25

There are many other factors and views to consider when looking to manage and optimize the end user experience in virtual desktops. For example, the above approach looks at a single user view of the problem. By leveraging your profile groups, you can simplify the approach on large scale implementation. Further, by elevating the above analysis from single users to groups, you can quickly determine the virtualization strategy and priorities at the corporate level. For example, what departments are prime candidates for virtualization? And related, which ones will deliver the strongest ROI/TCO?

Sizing the Host Memory

Now that you've established virtual profile and user groupings, it's time to consider the physical RAM required to support those groups. This step in the process begins with an estimation of the number of desktop profiles and the breakdown by users or user groups within each profile. Other considerations, such as resource variability by time, host hypervisor platform, video resolution and number of monitors all contribute to this exercise—prior to a final determination of sizing the host memory.

The first step in host memory sizing has little to no impact on the actual host sizing. Related, average memory usage per user or machine is also irrelevant; we turn first to the baseline CPU and memory requirements across all candidate users. The only exception here is to establish the maximum number of VMs per host, based on hypervisor settings and allocated memory.

Note that variability of user resource requirements over time is also important to consider. When presented as a timeline, as noted in the figure below, the combined metrics over time will show you how much resource would be consumed if all desktops were running in a shared environment. This is a great way to simulate the total infrastructure resource required.

In this example, you will note the peak combined RAM used is almost 3000MB (3GB) per machine. To support 400 VMs, this would represent a minimum RAM pool of 1.2TB; however, the actual RAM pool requirements will likely be higher based on other factors.

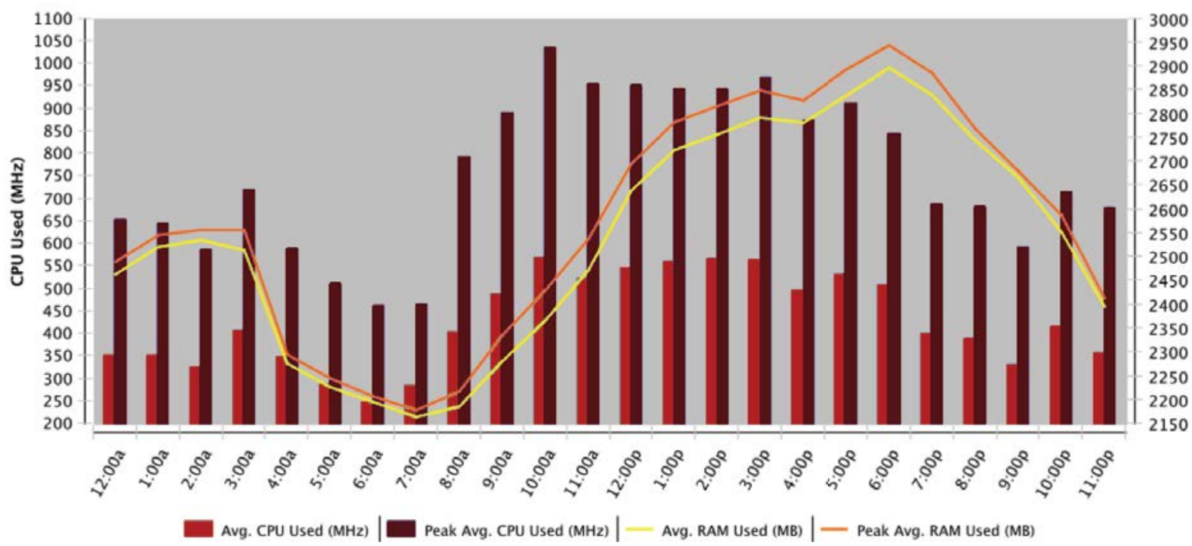

Based on the information observed, you can first establish the optimal host configuration by comparing CPU and memory usage to determine the number of servers required. Other adjustments that must be considered include:

- Adding memory overhead based on the host hypervisor platform
- Reduce memory requirement due to ballooning and page sharing (according to best practices)
- Adjust memory allocations based on remote display protocol (for example PCoIP vs. ICA/HDX)

PCoIP Client Display Overhead					
Display Resolution Standard	Width, in Pixels	Height, in Pixels	1-Monitor Overhead	2-Monitor Overhead	4-Monitor Overhead
VGA	640	480	2.34MB	4.69MB	9.38MB
SVGA	800	600	3.66MB	7.32MB	14.65MB
720p	1280	720	7.03MB	14.65MB	28.13MB
UXGA	1600	1200	14.65MB	29.30MB	58.59MB
1080p	1920	1080	15.82MB	31.64MB	63.28MB
WUXGA	1920	1200	17.58MB	35.16MB	70.31MB
QXGA	2048	1536	24.00MB	48.00MB	96.00MB
WQXGA	2560	1600	31.25MB	62.50MB	125.00MB

When you take all of the above into consideration, you will note that the actual allocation of virtual RAM to each VM does not change. That is, you do not need to allocate 1GB of RAM for applications and another 31MB for dual 1080p monitors. Instead, consider the overhead RAM when calculating the total physical RAM required for each server host. You simply add the guest VM RAM requirements to the overhead RAM requirements and multiply by the total number of virtual machines. By way of example, note how screen resolution and number of monitors have an impact on required memory.

Another aspect to consider when sizing host server memory is how the application landscape might change when migrating to a virtual environment. While impossible to predict each unique environment, we recommend a "balance sheet" approach to estimating application resource requirements. Take some time to identify and compare incoming/outgoing applications in an effort to estimate a possible increase or decrease in required resources for the desktop. Examples of common applications changes include:

- Outgoing: hardware monitoring (SMS), antivirus, software distribution, remote management
- Incoming: virtualization agent, application delivery, remote display protocol, profile management

Planning Ahead and Other Considerations

We strongly recommend including additional memory and CPU to the estimated requirements for both VMs and server hosts to account for future changes. Common changes that require additional resources include:

- Growth—most organizations include enough growth room to absorb the capacity equivalent to one or two host servers in the initial server host design. This is to account for easier scaling as well as for fault tolerance.
- Disaster recovery and failover—in addition to growth, planned infrastructure to support hardware failures is critical when moving desktop workloads from the PC to a centrally-managed data center workload.
- Management tool overhead—while some assume that new management tools will replace old management tools in a virtualized end user environment, we have observed an overall increase in the resources consumed by tools such as virtualization agents, profile management, etc.
- Operating system changes—moves from Windows XP to Windows 7, or changes in the CPU instruction set—like a move from 32 to 64-bit—can have a profound effect on virtual resources.

As noted earlier, other variables that must be considered include application delivery and your desktop virtualization approach (user sessions vs. virtual desktop, persistent vs. non-persistent). Each of these topics are well beyond the scope of this paper, but must be considered when looking at memory requirements for server host sizing. We will next turn our attention to storage considerations in the virtual desktop environment.

Storage Considerations for the Virtual Desktop Environment

In similar regard to the challenges and considerations for sizing memory for the virtual machine client and host, storage considerations to support virtual workloads can also be complex—and not to mention expensive, if not architected with care and attention to detail. In our experience, sizing storage appropriately for VDI is one of the most challenging aspects of building an optimal architecture. To be clear, the challenge is not so much "sizing" disk capacity, or even throughput—the biggest challenge is understanding storage input and output operations per second (IOPs). Broadly, this speaks to how quickly the system can read and write information to disk.

A common misconception of early adopters is to assume that desktops have, in general, less demanding disk workloads than servers. And while that is true when thinking in terms of throughput, the desktop IOPs requirements are far more complex. User IOPs are highly unpredictable. Most desktops generate random disk reads and writes, as opposed to sequential. Related, that pattern of use varies according to multiple factors, including: time of day, day of the month, application access, operating system events and type of user. Fitting the correct type of disk, RAID level and performance characteristics should be approached carefully. In short, you should never size your virtual desktop storage as you would for traditional server workloads.

Not only are desktop storage resource requirements unpredictable, they also can vary wildly. Specifically, this wild volatility between average and peak disk IOPs requirements can have far-reaching consequences to end user experience in the virtual desktop environment. In most cases, IOPs performance issues can be categorized as follows:

- The boot storm—easily addressed by provisioning VM ahead of login storm.
- Login storm—cannot be avoided. Requires a "flexible" storage architecture to absorb peaks without oversizing to the detriment of ROI.
- Application and user-generated peaks—can be eliminated, but require monitoring to detect and address.

One basic rule to mitigate storage-related performance issues is to avoid scheduled events. The more you can randomize events that tie to your storage infrastructure, the better you can optimize your environment. Of course this is more difficult to identify if you haven't measured these attributes and can plan accordingly.

Not Just About the IOPs

While we've highlighted storage IOPs as critical (and they are!), there are other attributes that should be considered when architecting the optimal storage environment to support virtual desktop workloads. In addition to IOPs, it's also important to consider attributes like read/write ratio, storage block size and how the desktop virtualization approach (persistent vs. non-persistent) will affect your implementation.

Related, disk capacity and leveraging the right storage technology for the right use case is an important consideration in balancing the user experience, overall project cost and maintaining flexibility. For example, desktop users are a tricky bunch. They typically find a way to save information wherever policy allows. In the world VDI, a more popular approach is one where you steer or redirect users (and the supporting storage functions) to the storage best suited to the required function. Like sizing host memory, sizing and tiering your storage resources should begin with a thoughtful approach. The figure above depicts average, peak average and peak read and write IOPs from our Stratusphere FIT Assessment product. Note the extreme performance characteristics observed, especially the gap between average and peak.

Metrics	Allocated	Average Used	Peak Average Used	Peak Used
Read IOPs		447.38 IOPs	943.9 IOPs	6157.41 IOPs
Write IOPs		175.88 IOPs	318.29 IOPs	2105.98 IOPs
Read Rate		2.27 GBps	6.19 GBps	55.67 GBps
Write Rate		1.59 GBps	5.06 GBps	55.64 GBps

Storage and the Right Tier for the Job

As we've noted, a metrics-driven approach is the optimal way to understand the various storage requirements within the overall supporting virtual desktop architecture. That approach should begin with a performance-based capacity planning exercise. Likewise, a similar metrics-driven approach should be employed when auditing an existing VDI environment.

And whether you're building from scratch, leveraging an existing investment in shared storage or optimizing an existing VDI architecture, the good news is there are many options to consider. For example, there are buffering and caching technologies available to augment existing storage solutions. The goal of this type of solution is straightforward: provide extra horsepower—by way of IOPs—to handle storage-related resource peaks.

You should be careful to compare these solutions as the market is currently crowded, and all solutions are not created equal. Be sure to understand the need and choose a technology and solution that provides the particular required storage attributes.

Building and Tiering

For those that are purpose-building a new storage architecture to support VDI, it is critical to understand the nature of the different VDI data types and how each contributes to your project. Overall the challenge is to keep capacity and cost under control, while delivering the horsepower where it's needed. For this reason we strongly recommend a tiered approach to architecting your storage. With virtualization, it is important to begin to think about VDI data as a few different types: operating system data, application data, and user and profile data. Not surprisingly, each of these VDI data types has a very different storage performance requirement.

- Operating system data—this type of data does not grow by very much over time. It is often very similar across desktops, and requires fast performance for the best user experience.
- Application data—application data storage requirements fluctuate moderately as applications are added, removed and updated. The core set of applications are typically shared across the majority of users; however, less-used application tend to align to specific user groups (or departments) and will vary based on user profile.
- User and profile data—this type of data is unique for each user. It is ever-constantly changing and the most critical data type in this exercise.

As we've already stated, the storage approach you should employ must tie back to the specific implementation goal—whether that be to support the operating system, applications or user and profile. In some cases the approach should be centralized to best take advantage of the similarities of the data. And in other cases, the data should be aligned to the user or profiles.

Earlier in this paper we discussed the desktop profile and grouping used based upon similarities in their resource consumption. This very same approach should be employed when looking at storage. By way of example, operating system data tends to be relatively static and would benefit from a centralized approach. We have found that standardizing the operating system within the master or golden VDI image, and placing this on your top-most tier of storage is optimal.

When looking at applications, the recommended approach is a more balanced one. For example, loading the core applications (those used by 80 percent or more of the user base) on the master or golden image will yield the greatest balance between performance and manageability. Lesser-used applications—including those used by single users, smaller groups and departments—should be segregated from the golden image, but still accessible and dynamically allocated to the virtual desktop session. The key to application-centric data is to align the applications to each user, not to the VM.

User and profile data is also critical to the overall performance and manageability of your end user environment. While outside the scope of this paper, it is important to consider the desktop virtualization approach (user sessions vs. virtual desktop, persistent vs. non-persistent), as this will have an effect on architecture and storage needs. Broadly, it is our recommendation that user and profile data should be completely segregated from VM data and stored in a user-based directory.

The Bottom Line

Architecting storage to support VDI requires that you understand how all the component elements will evolve and change over time—both in terms of the capacity and performance required. Remember these resource requirements will vary based on VDI data type. With respect to our findings here, it should be noted that the following storage characteristics are typical:

- Capacity tends to grow at a predictable rate
- Read and Write performance tends to be inconsistent
- IOPs can be highly volatile
- High IOPs/capacity ratio (as compared to server workloads)

For these reasons, a typical "cabinet" or shared storage approach may not be optimal when designing the infrastructure to support virtual desktop workloads. This is due to the interaction of IOPs and the number of spindles (number of discrete spinning drives) required to meet specific performance requirements. For this reason, a solution that includes some type of performance "buffer" is more favorable.

Building A New Storage Architecture: Solid State Disks

When it comes to VDI, all disk is not created equal, In fact, solid state disk (SSD) aren't spinning disks at all; hence their use in VDI has become commonplace. Traditional hard drives have difficulty responding to the random access that is prevalent in virtual desktop workloads. In smaller VDI environments this challenge can be mitigated by throwing more traditional hard disks at the problem. As you grow the number of VMs in the environment, the more likely the benefit you will realize with SSD-based storage.

We typically recommend SSDs to help mitigate VDI boot storms, for example. In some cases it can be as simple as placing collections of required files on an SSD tier—as a means to create a caching layer. In other cases, the need may be to provide higher levels of storage IOPs to deal with instances of peak or extended high demand.

As noted prior, SSDs play a critical role in supporting a tiered storage approach. Breaking down different data types (operating system, applications, user and profile) and making a thoughtful determination based upon factors like data change rate, user resources tied to the storage, etc. should be considered. At a high level, we would make the following generalizations:

- SSD—for base image. Excellent for optimizing performance and handling boot/login storms.
- HD—for user and profile data as well as non-critical applications.
- SSD or HD—for more critical applications, as budget allows.

Regardless of the drive technology, you must quantify your use cases. The steps are similar to those discussed to support memory sizing:

- Gather assessment data
- Calculate average and peak average IOPs per client
- Identify peak single time average across user base
- Employ a strategy that delivers more IOPs than peak single time average

Calculating average and peak average IOPs per client as a means to ensure end user experience is an important first step in any VDI project. Our assessment tool, Stratusphere FIT, offers an Assessment Findings report that provides aggregated read and write IOPs for both average and peak values. Understanding these attributes are critical to properly sizing and tiering your storage environment to support virtual desktop workloads.

Using Assessment Data to Calculate Resource Requirements

By now you've heard us make the point that storage is the most overlooked, but one of the most critical components of a successful VDI implementation. We have seen our share of shotgun approaches to storage to support VDI. And while short term gains and remediation are sometimes possible, we've also seen our fair share of projects where an overbuilding exercise creates an environment that will never receive a positive ROI. We've also made a similar point regarding host server memory; both server memory and storage are critically important to virtual desktop workloads.

We'd like to conclude by sharing a few recommendations with respect to the characteristics you should strive to provide in your VDI architecture:

- Host servers should be built with a CPU/RAM ratio adapted to your user population. Your ability to recognize the strongest ROI and TCO characteristics on the host server platform are cast in these resource elements. You will also eliminate the risk of performance bottlenecks, resulting in poor user experience.

- Identify the appropriate user groupings based on your unique user population. As noted, it would be the norm for all end user virtualization efforts to have two desktop profile groups, at a minimum. Design each pool to ensure your user groups get what they need without building to the highest common denominator (higher user experience at a high resource expense) or lowest (poor user experience at a minimal cost).
- Application, user profile and data delivery should focus on the user, not the machine. Related, architectural decisions like "follow me" personalization will offer profound flexibility, but also help you mitigate risk in your environment.
- A multi-tiered storage design will yield maximum flexibility and allow you to contain costs with respect to over-building. Purchase only what you need (performance and/or capacity) without being locked into a single storage configuration. This must include a performance or buffer component (flash, SSD, etc.) to accommodate peak resource needs.

Start your project with a thoughtful assessment strategy and data capturing process. Just as every human being is different, every desktop workload is different as well. Measure as many of your total desktop users as possible—ideally collecting assessment information from every user in your organization. The path to success is observing solid metrics-based data before you make a commitment to architectural components. And with the right VDI architecture, you can ensure you are balancing the user experience, manageability, flexibility and ultimately ROI and TCO attributes.