



Technical Report

Guidelines for Virtual Desktop Storage Profiling and Sizing

Angela Ge, Keith Aasen, NetApp
March 2011 | TR-3902

TABLE OF CONTENTS

- 1 INTRODUCTION TO STORAGE SIZING GUIDELINES FOR HOSTED VIRTUAL DESKTOP ENVIRONMENT 3**
- 2 HVD STORAGE SIZING FUNDAMENTALS 3**
 - 2.1 KEY CONSIDERATIONS IN HVD STORAGE3
 - 2.2 BEFORE YOU SIZE.....4
 - 2.3 HVD CAPACITY REQUIREMENTS4
 - 2.4 HVD PERFORMANCE REQUIREMENTS.....5
 - 2.5 READ IOPS6
 - 2.6 WRITE IOPS7
- 3 SIZING PROCESS 7**
 - 3.1 STEP ONE: GATHER INPUTS8
 - 3.2 STEP TWO: RUN NETAPP VDI STORAGE SIZER17
 - 3.3 STEP THREE: MINIMIZE STORAGE CONSUMPTION17
- 4 ACKNOWLEDGEMENTS 19**
- 5 FEEDBACK..... 20**
- 6 VERSION HISTORY..... 20**
- 7 REFERENCES 20**

LIST OF TABLES

- Table 1) Example that illustrates the math. 5
- Table 2) Continuing our example: How many drives are required to serve read IOPS? 6
- Table 3) Now let’s look at the storage needed to meet the write performance requirement. 7
- Table 4) Peak IOPS calculation..... 11

LIST OF FIGURES

- Figure 1) Machine profile. 9
- Figure 2) Profile characteristics. 10
- Figure 3) Application inventory. 13
- Figure 4) Stratusphere Fit: Reports screen. 14
- Figure 5) Stratusphere Fit: report library..... 15
- Figure 6) Stratusphere Fit: create report. 16
- Figure 7) Stratusphere Fit: run report. 16
- Figure 8) File-level misalignment..... 18
- Figure 9) Proper alignment of guest file system. 18

1 INTRODUCTION TO STORAGE SIZING GUIDELINES FOR HOSTED VIRTUAL DESKTOP ENVIRONMENT

This document provides guidelines for NetApp® storage sizing in the hosted virtual desktop (HVD) environment. HVD is also sometimes referred to as virtual desktop infrastructure (VDI). This document addresses HVD sizing fundamentals, outlines a recommended process for proper storage sizing, discusses common culprits of storage overconsumption, and provides best practices for reducing storage consumption. This document is applicable to VMware®, Citrix, Microsoft®, Quest, and DeskTone HVD environments, as well as other HVD vendors. The specific examples provided in this document are based on assumptions about customer environments and are for illustrative purposes only. Customer environments vary greatly; therefore, it's important to assess your customer's unique environment before sizing.

The intended audience for this document is NetApp sales engineers (SEs), NetApp Professional Services (PS), and channel partners with a good basic knowledge of storage and HVD technologies.

Important

This document is not a substitute for NetApp sizing tools. Always consult the appropriate sizing tool and/or a NetApp sizing expert for the final configuration. The VDI storage sizing tool is available for partners at <https://fieldportal.netapp.com/sites/sizers/index.html> and for NetApp employees at <http://perf-build1.eng.netapp.com/VDISizerApp/Display>.

2 HVD STORAGE SIZING FUNDAMENTALS

2.1 KEY CONSIDERATIONS IN HVD STORAGE

As more enterprise desktops move from distributed computing environments to HVD in centralized data centers, it is imperative to design a robust and cost-effective storage back end. Two requirements must be taken into consideration: (1) capacity and (2) performance. They directly affect storage sizing and, ultimately, user experience and cost per desktop.

Storage sizing for HVD can be complex. Undersizing results in an inadequate amount of storage to serve the performance requirement of the desktops and storage bottlenecks during steady-state usage and various peak times, such as boot storms, logon storms, resync storms, virus scans, and patch updates. Oversizing is almost equally bad and leads to expensive cost per desktop projections and stalled HVD projects. The goal is to size just enough storage to meet both the capacity and performance demands of the HVD environment. Storage cost is one of the key barriers to HVD growth; getting the sizing right is critical to project success.

The NetApp VDI storage sizer, with properly gathered inputs, can size a configuration that meets both capacity and performance demands. The following sections illustrate, by using a simple example, what goes on inside the sizer.

Note: The example shown in this report is for illustrative purposes only. Do not attempt to size your virtual environment by using the assumptions and results in the example. Instead, use the NetApp VDI sizer. The sizer has a sophisticated back-end engine and considers many variables in addition to what is included in our example. Therefore, the output might not conform exactly to the example. All sizing examples given in this document assume that Network File System (NFS) is used. Sizing might vary on other protocols. Consult the VDI sizer or a NetApp expert for more details.

Important

Consulting with the desktop team as soon as possible is critical.

2.2 BEFORE YOU SIZE

Capacity and performance requirements are driven by characteristics of the HVD environment. Therefore, learn as much as you can about your environment before you start sizing. The following list contains key questions that you should ask:

- Who will use these HVD desktops? Examples include general-purpose knowledge worker, data entry, kiosk, developers, outsourced workers, and so forth. What are their customization requirements? Examples include full customization (that is, the ability to download applications), some customization (the ability to change user settings but not to download applications), and none at all (that is, desktop is fully locked down and no changes are allowed). This information can help you decide whether the virtual desktops need to be stateful or stateless, as well as figure out how to handle patching and updating. Both affect storage sizing and deployment.
- What are the patterns of desktop use (random, shift work, multiple shifts)? This information helps determine the degree of concurrency in desktop usage and desktop login. Higher concurrency requirements typically lead to higher storage consumption.
- What notable applications are being used today? Will all of these applications be migrated to the HVD environment? Is it possible to drop some applications or to consider alternative delivery methods (for example, application streaming) for some applications? Fewer applications installed in the HVD environment lead to lower performance requirements and storage savings.
- How will user data be handled? Will it reside in the virtual machines (VMs), or will it be redirected to a CIFS share? Keep in mind that separating user data from virtual desktops reduces storage consumption. User data typically has much lower performance requirements than virtual desktop data and can therefore be put on less costly SATA drives.
- What is the growth or deployment plan for the desktops? That is, how many are needed today, how quickly do you expect to grow, and how many more desktops will be needed? This information helps determine the appropriate platform.

Answers to these and other questions will help you determine the inputs needed for capacity and performance sizing. This list of questions, which is by no means complete, is meant to get you started on the journey to thoroughly understanding the customer environment. The remainder of this document discusses additional details that you must consider in the sizing process.

2.3 HVD CAPACITY REQUIREMENTS

Capacity estimation is relatively straightforward. Be sure to account for all capacity needs, including C: drives and user home directories. The typical capacity range for virtual desktops could be anywhere from 5GB to 30GB. When user data is consolidated and redirected to a file share and applications are isolated and streamed, the capacity requirement per desktop tends to be lower. Factor in savings with multiple levels of storage efficiency technologies such as thin provisioning, deduplication, and FlexClone[®] volumes. Technologies provided by such solutions as VMware View Composer linked clones or the Citrix XenDesktop[®] and Provisioning Server can also be used to provide storage efficiency.

For virtual desktops (not including user home directories), we typically see over 95% savings in capacity after the storage efficiency technologies are applied. However, reducing the capacity requirement of virtual desktops is the first step, because the storage required to meet the performance requirements often exceeds that of the now reduced capacity requirement.

Table 1 illustrates the math. For simplicity, we excluded user data and a storage overhead calculation.

Table 1) Example that illustrates the math.

Item	Amount
GB per desktop	20GB
Number of desktops	1,000
Initial capacity	20TB
Dedupe savings	95%
Postdedupe capacity	1TB

Unlike virtual desktops, user home directory data typically has a low performance requirement. Therefore, storage sizing for home directories is very often driven by capacity requirements alone. Customers can expect to see a storage savings of somewhere from 35% to 50% on average when using NetApp deduplication on home directories. Combined with the new compression technology available in Data ONTAP 8.0.1, storage savings on home directories can reach 60% to 70%. Keep in mind, however, that these are just averages. Environments differ from one another; therefore, the savings might vary. NetApp provides a separate Common Internet File System protocol (CIFS) sizing tool for sizing home directory storage usage. The VDI storage sizing tool is available for partners at <https://fieldportal.netapp.com/sites/sizers/index.html> and for NetApp employees at <http://perf-build1.eng.netapp.com/CIFSSizer/CIFSizer.jsp>.

2.4 HVD PERFORMANCE REQUIREMENTS

A proper estimation of performance requirements is crucial to HVD storage sizing, and it is typically the most difficult part of the project. Most customers do not have a good understanding of the performance requirements of their users. Therefore, NetApp highly recommends the use of a third-party assessment tool, such as Stratusphere Fit™ from Liquidware Labs. To make this tool easier to use, NetApp collaborated with Liquidware Labs to develop a NetApp custom report template in Stratusphere, which aggregates all the necessary sizing inputs. [Section 3](#) discusses Stratusphere and the NetApp custom report in detail.

IOPS is the standard unit of measurement for input/output (I/O) operations per second. This measurement is a performance-based measurement and is usually written as:

- **Total IOPS.** Average number of I/O operations per second
- **Read IOPS.** Average number of read I/O operations per second
- **Write IOPS.** Average number of write I/O operations per second

The read I/O percentage is the ratio of read IOPS to total IOPS. The write I/O percentage is the ratio of write IOPS to total IOPS.

HVD workloads tend to be “bursty”: IOPS required during certain operations such as booting, login, and antivirus (AV) scans are significantly higher than for normal steady-state operations during most of the workday.

The key performance drivers of sizing a normal steady state are IOPS per desktop and the read/write ratio. IOPS per desktop vary widely, depending on the user profile. Estimating IOPS per desktop requires a clear understanding of the user base and how they use their desktops, assessment tools such as Stratusphere, and running a proof of concept (POC).

The read/write ratio also varies from environment to environment. The HVD workloads, however, tend to be more write-heavy at steady state. We have seen read/write ratios of 25/75 and 20/80 with many customers, and sometimes even 10/90. Another consideration is whether the IOPS are random or sequential. Most HVD IOPS are random in nature.

The same set of key performance drivers needs to be considered for peak times. During peak times, IOPS required per desktop will be higher, and the read/write ratio might be different.

NetApp technologies such as Flash Cache, deduplication, and WAFL[®] (Write Anywhere File Layout) enable performance requirements to be met with a minimum number of disk drives. The following sections take a separate look at reads and at writes to see how the various technologies apply.

2.5 READ IOPS

Serving read IOPS is straightforward on a traditional third-party system (one read IOP results in one disk IOP). The NetApp intelligent caching technology significantly reduces the number of disk drives needed to sustain read IOPS and breaks the one-read IOP to one-disk IOP limitation. In typical HVD environments, we often see a 75% disk reduction for normal steady-state read IOPS and the ability to handle I/O bursts during peak conditions. The NetApp intelligent cache consists of three components:

- Native algorithm in Data ONTAP[®] 7.3.1 and beyond, which is available on all FAS platforms
- Flash Cache (formerly Performance Accelerator Module, or PAM): flash-based memory cache in FAS3X00 and FAS6X00 platforms that extends the native algorithm for greater scalability
- Deduplication: a highly deduplicated environment such as HVD means that all the frequently used data (for example, OS) has a small footprint and is cached in the memory for much faster retrieval

Continuing our example, Table 2 illustrates the amount of drives required to serve read IOPS on both a third-party and a NetApp system.

Note: These figures are for illustrative purposes only. Do not use them to size your environment.

Table 2) Continuing our example: How many drives are required to serve read IOPS?

Item	Amount
IOPS/user (normal hours)	8
IOPS/user (peak hours)	16
Read/write	50/50
Number of read IOPS required	$16 \times 1,000 \times 50\% = 8,000$
Number of IOPS per 15K 300GB SAS drive	220 (for illustrative purposes only)
Number of drives required for reads on third-party system	$8,000 / 220 = 36$
Percentage of read IOPS offloaded by Flash Cache	75%
Number of remaining read IOPS with Flash Cache	$8,000 \times (100\% - 75\%) = 2,000$
Number of drives required for reads on NetApp	$2,000 / 220 = 9$

Common Misconception

A common misconception is that Flash Cache enables SATA disks to be used in an HVD environment. Flash Cache offloads a significant amount of read IOPS. However, the remaining read and write IOPS are still more cost-efficiently served from SAS/FC disks.

2.6 WRITE IOPS

Similar to reads, HVD writes are typically random in nature. Random writes are the most expensive operation for almost all RAID systems. RAID achieves resiliency by striping data across multiple disks, and random writes mean that each write operation requires more than one disk operation. The ratio of client write to disk write depends on the RAID type. For example, in a RAID 5 configuration on a traditional storage array, each client write IOP requires up to four disk IOPS. A large write cache on traditional storage systems might help, because write coalescing of requests happens if the write cache is big enough, but it still requires at least two disk IOPS. In a RAID 10 configuration, each client write IOP requires two disk IOPS; however, RAID 10 requires twice the number of disks as RAID 5.

NetApp Data ONTAP and WAFL, standard in all NetApp systems, have been architected since 1993 to specifically optimize write IOPS for RAID-DP[®]. Data ONTAP coalesces multiple write IOPS and sends them to disk as a single IOP. This effectively turns random writes, which is the most expensive IOP for disks, into much less costly sequential writes. On a NetApp system, multiple client write IOPS correspond to one disk IOP. Compared to storage arrays with RAID 5 or RAID 10, which require at least twice the disk IOPS per client IOPS, write operations require significantly fewer disks on NetApp storage.

Table 3 shows the storage needed to meet the write performance on third-party storage and NetApp storage, using the same example.

Note: These figures are for illustrative purposes only. Do not use them to size your environment.

Table 3) Now let's look at the storage needed to meet the write performance requirement.

Item	Amount
IOPS/user (normal hours)	8
IOPS/user (peak hours)	16
Read/write	50 / 50
Number of write IOPS required	$16 \times 1,000 \times 50\% = 8,000$
Number of IOPS per 15K 300GB SAS drive	220 (for illustrative purposes)
Number of SAS drives required for writes on third-party storage	$8,000 / 220 = 36$
Number of SAS drives required for reads and writes on third-party storage	$36 + 36 = 72$
Number of SAS drives required on NetApp for reads and writes	$9 + 18 = 27$

Note: This is a conservative estimate of the maximum number of drives needed on NetApp. This assumes that WAFL coalesces two client write IOPS into one disk IOP. Actually, the WAFL algorithm is often more efficient and coalesces much more than two writes.

3 SIZING PROCESS

The storage sizing process can be broken down into three steps. Each step is discussed in more detail in the following sections:

- Step 1: Gather inputs
- Step 2: Run NetApp VDI sizer
- Step 3: Minimize storage consumption

3.1 STEP ONE: GATHER INPUTS

LIQUIDWARE LABS STRATUSPHERE FIT

NetApp is partnering with Liquidware Labs, a leading desktop transformation solutions company. The Liquidware Labs Stratusphere Fit leads the market in gathering data and metrics of all integral infrastructure needed for HVD, from the desktop, to the network, to the data center, and to the storage. The solution features dozens of built-in reports to help you plan for your next desktops. Stratusphere Fit identifies ideal users and desktops for HVD, assesses ideal applications for application virtualization, and contains features to help you size storage appropriately. Stratusphere Fit also features specific custom written reports for NetApp environments that greatly assist in assessing and sizing the storage needed for HVD environments.

For more information, including a free 10-day trial download of Stratusphere Fit, visit www.LiquidwareLabs.com.

NETAPP CUSTOM REPORT FOR STRATUSPHERE FIT

Although Stratusphere Fit can produce hundreds of performance metrics, determining which ones are essential for storage sizing and getting them in the format that can be consumed by the NetApp VDI sizer are not easy tasks. To make the process easier, NetApp has collaborated with Liquidware Labs to create a NetApp custom report template. By running the NetApp template in Stratusphere, you can generate a report that contains all the key variables needed for the sizing process. The parts of the sample report are discussed in this section and are shown in Figure 1, Figure 2, and Figure 3.

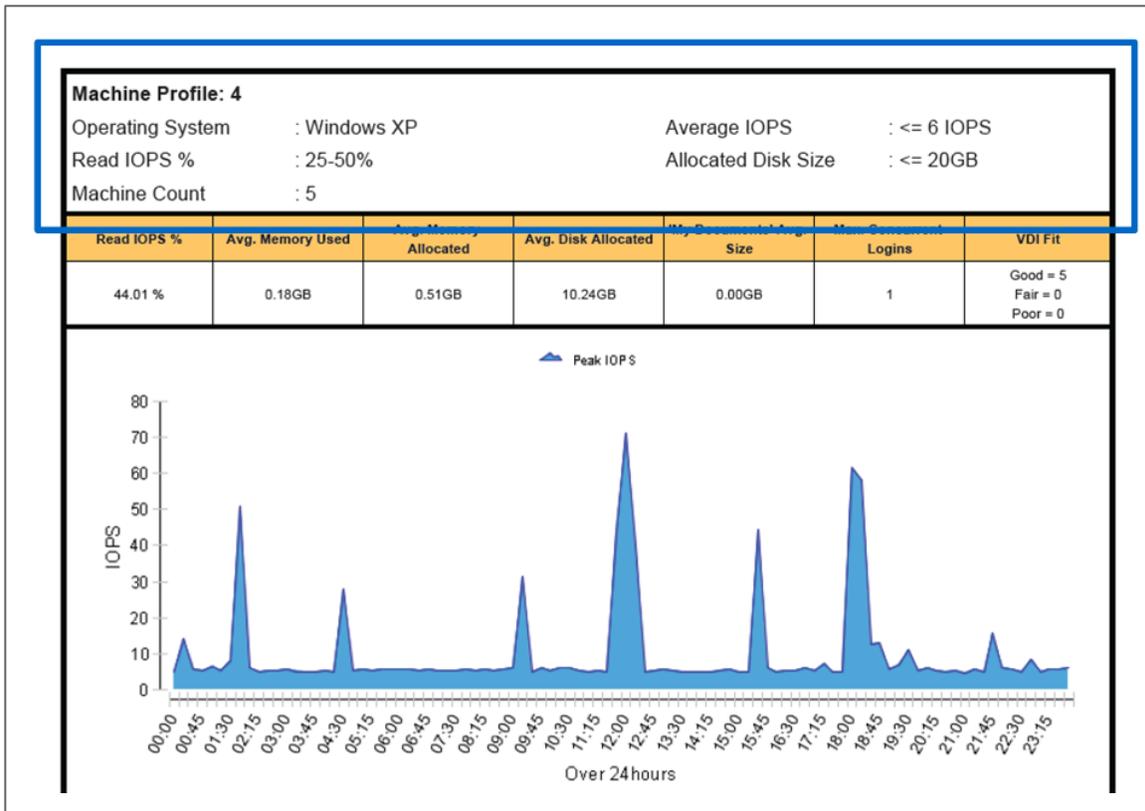
Note: The sample report is generated from a group of machines in the lab environment that are kept mostly idle; therefore, the report is not representative of an actual environment. IOPS from actual customer environments tend to be much higher. These sample metrics are for illustrative purposes only; therefore, do not use them to size your environment.

Machine Profile

Stratusphere can be run on a group of machines (either a physical PC or VM) over a desired period of time. The measurement period should be long enough to cover the typical set of activities performed by the user population in the environment. NetApp recommends running the tool over a 24-hour period at the very minimum and ideally over 30 days to capture all of the peaks. Each machine within the group might vary in usage characteristics. Machine profile is a methodology to group similar machines into a profile. Machine profiles are labeled numerically in incremental order. Each machine profile in the report corresponds to a user profile in the NetApp VDI sizer.

The machine profile section of the report is highlighted in Figure 1.

Figure 1) Machine profile.



The NetApp custom report template assigns machines into various machine profiles based on the following criteria and segmentation:

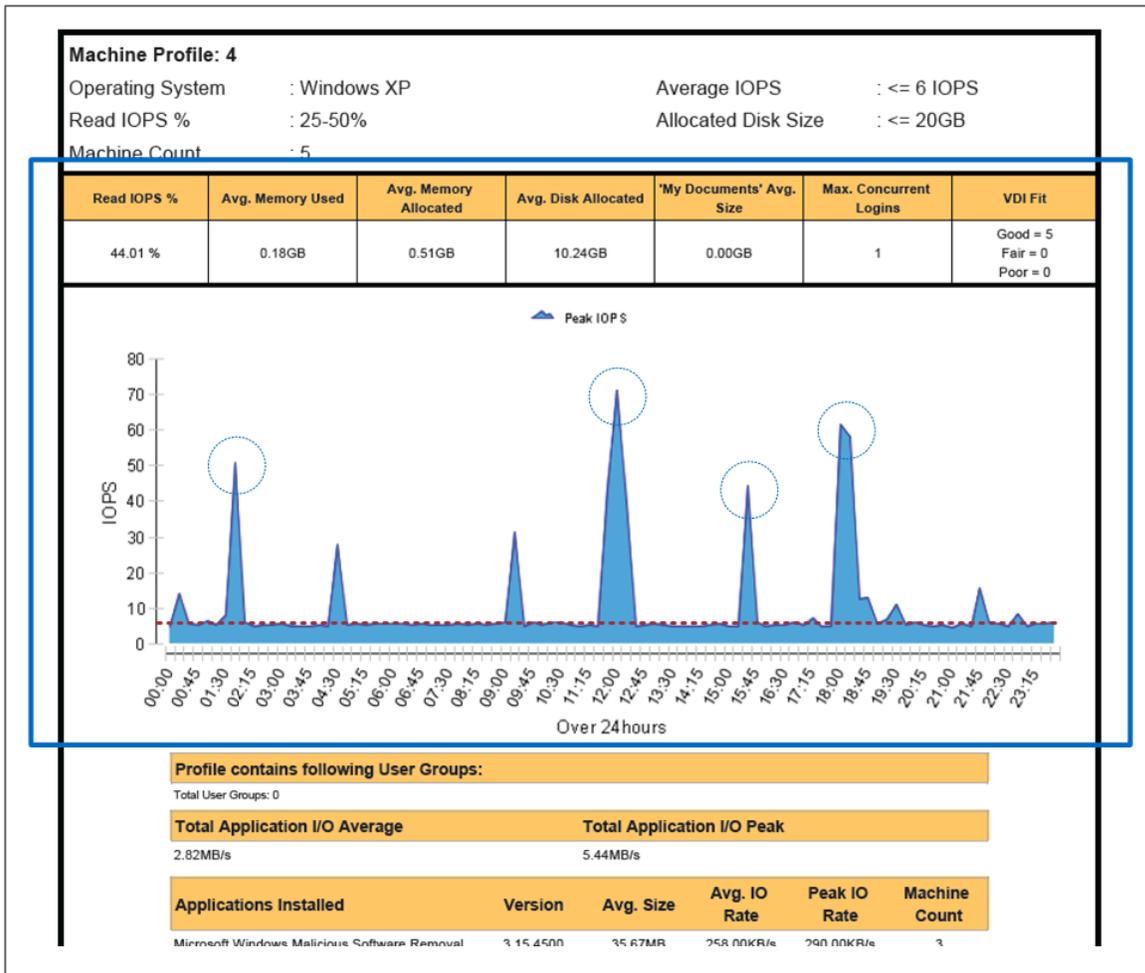
- Operating system (Windows® XP, Windows Vista®, Windows 7, and so forth)
- Read IOPS as a percentage of total IOPS (<25%, 25% to 50%, 50% to 75%, 75% to 100%)
- Average IOPS (<6, 6–12, 12–18, 18–24, 24+)
- Allocated disk size (<20GB, 20–40GB, 40–60GB, 60–80GB, 80GB+)

Machine count shows the total number of machines placed in a specific machine profile. The example shown in Figure 1 has five machines in this machine profile. All of these machines run the Windows XP operating system. Each machine has an average IOP of less than or equal to 6. The average read IOP ratio of each machine is from 25% to 50%. The allocated disk size on each machine in this group is less than or equal to 20GB.

Profile Characteristics

The profile characteristics portion of the report is highlighted in Figure 2.

Figure 2) Profile characteristics.



For each machine profile, the report shows the following metrics:

- **Read IOPS %.** The average read ratio of all the machines in the machine profile during the measurement period. Although this measurement can indicate how your users use their desktops pre-HVD, migrating these users to HVD might alter the read/write pattern due to the presence of connection brokers and changes to the OS and applications. Therefore, consider the impact of the migration on the read IOPS percentage before inputting this metric into the read percentage field in the NetApp VDI sizer. See the guidelines given in [section 2.4](#). For the most accurate measurement of the read and write percentage, conduct a POC.
 - **Avg. Memory Used.** The average memory working set used across all the machines in this machine profile.
 - **Avg. Memory Allocated.** The average memory working set allocated across all the machines in this machine profile. If the average memory allocated is much greater than the average used, you might consider lowering the memory allocation in your HVD environment. However, if you are more risk averse, you might want to leave a bigger margin. Input this value into “VM Memory Size (GB)” in the NetApp VDI sizer.
- Note:** If you are upgrading your OS (for example, from Windows XP to Windows 7), your memory allocation might be different. For Windows 7, we typically see an average memory allocation of 1.0–1.5GB per VM.

- **Average Disk Allocated.** The average capacity allocated across all the machines in this machine profile. This can serve as a reference point for the input to C: drive size (GB) in the NetApp VDI sizer. Keep in mind, as you move from a physical PC environment to a virtual desktop environment, that you might decide to implement a new strategy for hosting user data and delivery applications (for example, using application streaming). Therefore, the C: drive size you end up with might be very different from the average disk allocated metric from the profiling exercise.
- **“My Documents” Avg. Size.** The average gigabytes allocated to the My Documents folder across all the machines in this machine profile. This metric is a good indicator of how much user data your users typically store, and it might serve as a reference point for determining the user data quota to be stored in a file share when you move to the virtual desktop environment.
- **VDI Fit.** A proprietary score developed by Liquidware Labs to assess whether the virtual desktop is a good fit for this pool of users. Consult the Liquidware Labs Stratusphere manual for more details on this score.
- **Peak IOPS Chart.** The chart plots the maximum of the sum of all the machines’ IOPS during each 15-minute interval within the measured time period. IOPS measurements are taken every 5 minutes within the 15-minute interval. Table 4 shows an example of how the peak IOPS are calculated over a 15-minute interval and what the chart plots. If the machine profile group only has three machines, then during the 15-minute interval, three measurements are taken at times 5, 10, and 15. At time 5, machine 1 exhibits IOPS of 6, machine 2 has IOPS of 9, and machine 3 is idle and therefore has 0 IOPS. The sum of IOPS at time 5 for all of the machines is 15 IOPS. Using the same logic, the sum of IOPS measured at time 10 for all of the machines is 27 IOPS, and the sum of IOPS measured at time 15 is 20 IOPS. So the maximum sum for all of the machines is 27 IOPS. This sum is plotted on the chart for that particular 15-minute interval.

Table 4) Peak IOPS calculation.

	IOPS at Time 5	IOPS at Time 10	IOPS at Time 15
Machine 1	6 IOPS	12 IOPS	8 IOPS
Machine 2	9 IOPS	6 IOPS	2 IOPS
Machine 3	0 IOPS	9 IOPS	10 IOPS
Sum of peak IOPS	15 IOPS	27 IOPS	20 IOPS

The sample IOPS chart report has some interesting IOPS to look for. The first is the steady-state maximum IOPS across all the machines, which is represented by the dotted red line. For the majority of the 24-hour period, 6 IOPS is the maximum total IOPS across all of the five desktops in the profile at steady state. In addition, several peaks seem to occur, and some are larger than others. Some of the peaks are represented by blue circles in Figure 2. The major question is which IOPS amount do we use to size the storage: Should we use steady state IOPS or one of the peaks? Ask the following questions about the peaks:

- What event causes each of the peaks? Is it because the users are concurrently booting their desktops or logging onto their desktops?
- Are applications running on the desktops at the same time?
- Is this an antivirus scan or an update?

The answers to these questions help determine whether you can eliminate or mitigate the effect of these events on IOPS once these users are migrated to the HVD environment. If the cause is booting, could that be staggered in the HVD environment? If the cause is antivirus scanning or a definition update, could antivirus software that is specifically designed for HVD be used? If the cause is a “blacklist application,” there might be a better way to deliver them (a more detailed discussion of blacklist applications occurs later in this document).

Each peak above steady state needs to be addressed separately to see if the event that is causing the peak can be eliminated in the new HVD environment. If all peaks can be eliminated, then use the IOPS at the steady state. You might not be able to eliminate some peaks; therefore, size for the IOPS at the highest peak that you cannot eliminate.

For example, after investigating the peaks, you determine that the two highest peaks can be eliminated in the HVD environment. They are both caused by machines booting at the same time. Therefore, they can be eliminated by staggering the machine boots in the HVD environment. But the third-highest peak is caused by an unavoidable event. Therefore, size for this peak, which has a total sum IOPS of approximately 50 across all 5 desktops. The IOPS input into the NetApp VDI sizer is in the form of average IOPS. So divide the total IOPS by the number of machines in the profile. In this case, input 10 IOPS (or 50 IOPS / 5 machines) into the Avg. IOPS field in the NetApp VDI sizer. Certain peak events, such as booting and logging in, might have different IOPS requirements in pre-HVD and HVD environments because of the presence of the connection broker or other management changes. Therefore, consider the effect if you must size for these peak events. POC is typically the best way to get an accurate measurement.

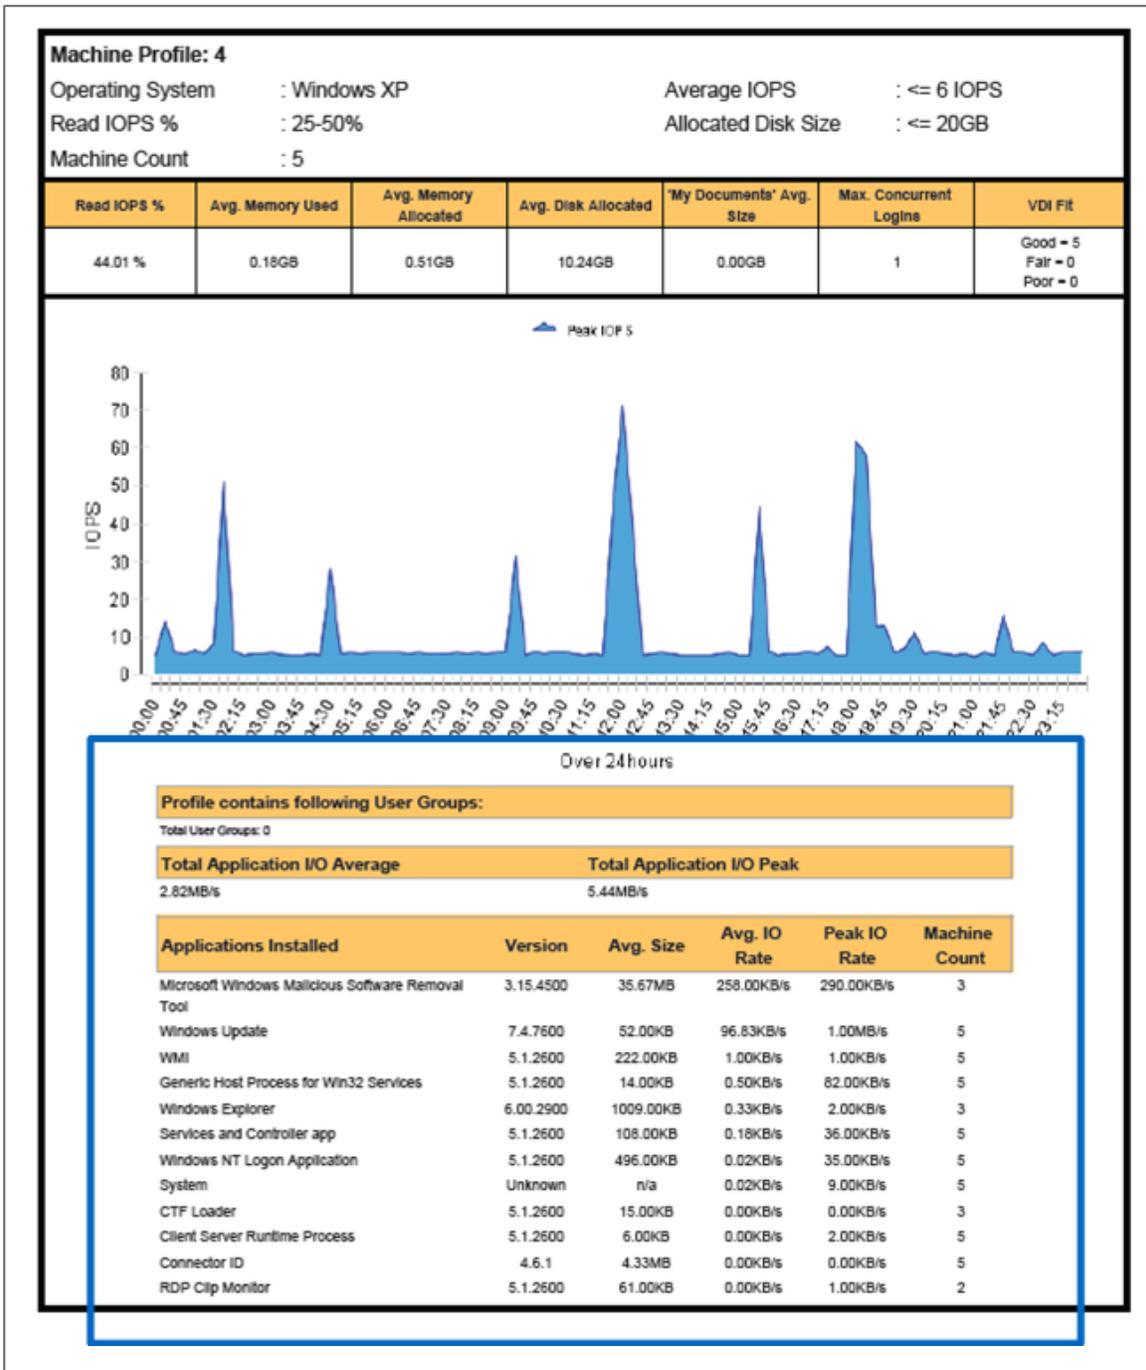
NetApp technologies such as Flash Cache and WAFL are used in the HVD environment to optimize performance. In many cases, the peak sizing might be similar or might only differ slightly due to the effect of these optimization technologies.

- **Max. Concurrent Logins.** The maximum number of users logging into their machines during any 10-minute window over the measured time period. In the example given, three users logged into their machines within 10 minutes of one another, which represented the maximum number of concurrent logins during the 24-hour measured time period. In an HVD environment, concurrent logins by a large percentage of users tend to cause I/O bottlenecks. This metric can help you determine the severity of the impact. For example, you might have 500 users, and roughly 300 of them might log into their desktops during any 10-minute window (for example, a shift change at a call center). Through your profiling exercise, you have determined a steady-state IOPS of 10 per user. Therefore, you size for 5,000 total IOPS. Then, you run a POC to measure login and discover that each login consumes 20 IOPS; therefore, you need a total of 6,000 IOPS for mass login. In this case, you might need to rerun the sizer.

Application Inventory

For each machine profile, in addition to the average metrics of the machine profile, the report also shows application-level information. This valuable information can be used to ascertain why a group of machines might have a higher IOPS requirement, as well as being used to lower IOPS and therefore storage cost per desktop. The application inventory section of the report is highlighted in Figure 3.

Figure 3) Application inventory.



This section of the report displays the following information:

- Application Installed and Version.** A list of all the applications (and versions) that are installed across the group of machines in the machine profile. A quick review will likely uncover unsanctioned applications on the users' desktops as well as help determine which applications should be migrated to the virtual desktop environment. [Section 3.3](#) covers the types of applications that are likely candidates for reevaluation.

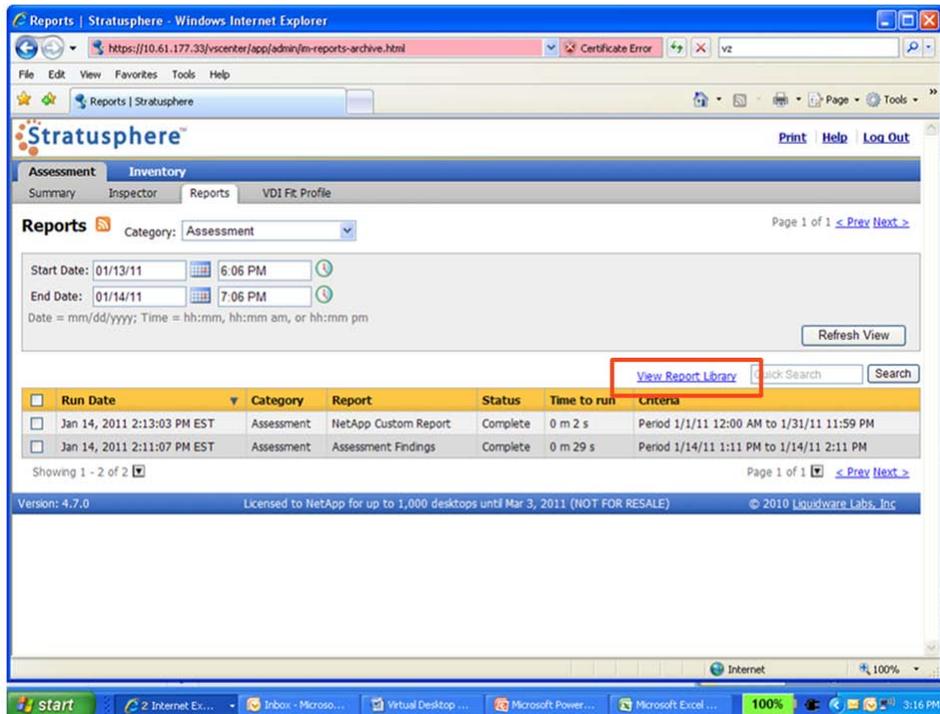
- **Avg. Size.** The average size of the particular application across all the machines in the machine profile group.
- **Avg. I/O Rate and Peak I/O Rate.** The throughput required by each application on average and at the peak during the measurement period. These two metrics can point to the applications that are “I/O hungry” and can lead to an investigation into alternative delivery methods for these applications in the new virtual desktop environment.
- **Machine Count.** The number of machines in the machine profile group that are running this particular application.

ACCESSING THE NETAPP CUSTOM REPORT FROM STRATOSPHERE FIT

To run the NetApp custom report, you must first download and install Stratusphere Fit and then save the NetApp custom report template file in a designated directory of your choice. Complete the following steps:

1. From the Assessment tab of the main screen, click the Reports subtab. Then click View Report Library.

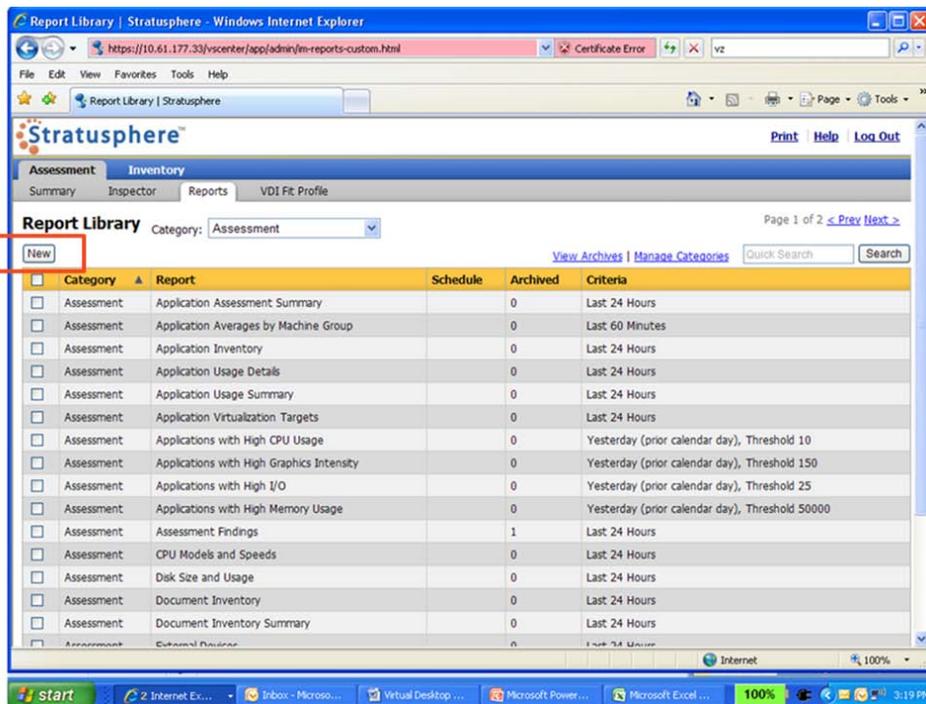
Figure 4) Stratusphere Fit: Reports screen.



A list of all of the report templates displays in the report library.

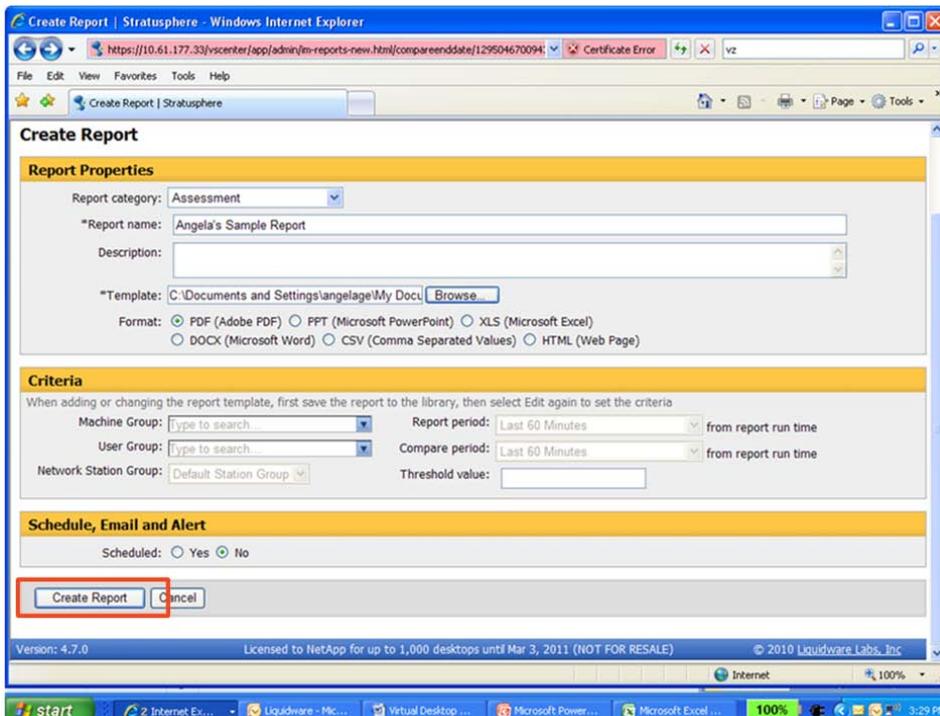
2. To create a new report, click New to open a new report wizard.

Figure 5) Stratusphere Fit: report library.



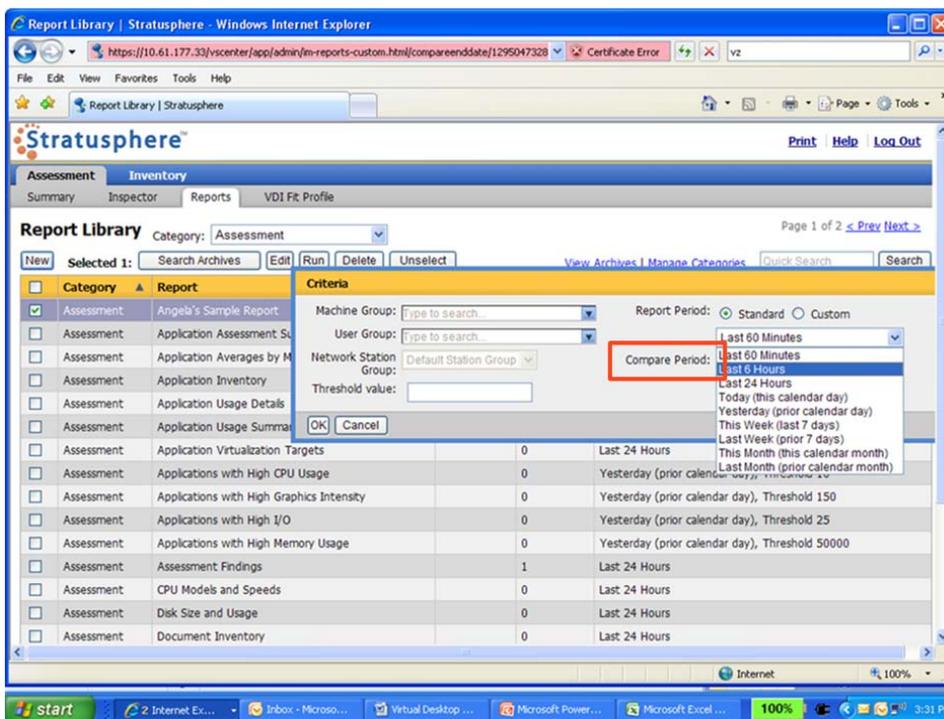
3. Select Assessment for the Report category field.
4. Create a new name for the report.
5. Load the NetApp custom report template you had previously saved. The report can be generated in .pdf, .ppt, .xls, or other formats.
6. Select your desired format.
7. Click Create Report to generate the report template.

Figure 6) Stratusphere Fit: create report.



8. Select the report template you just created in the report library and click Run. This opens a criteria wizard for the run.
9. Select your desired assessment period and click OK.

Figure 7) Stratusphere Fit: run report.



10. Click View Later in Archive if you want to navigate away from this screen. Otherwise, when the report is finished, click View to read the report.

3.2 STEP TWO: RUN NETAPP VDI STORAGE SIZER

After you have generated the NetApp custom report in Stratusphere Fit, the next step is to enter the inputs into the NetApp VDI sizer to generate a preliminary sizing and configuration. The sizer is available to all NetApp employees at <http://perf-build1.eng.netapp.com/VDISizerApp/Display> and to channel partners at <https://fieldportal.netapp.com/portals/default.asp>. For more details on how to use the sizer, consult the sizer FAQ and your NetApp SE.

3.3 STEP THREE: MINIMIZE STORAGE CONSUMPTION

In the previous section, we discussed one way of minimizing storage consumption by examining the throughput consumptions of the applications on desktops and rethinking ways of delivering these applications. This section contains a more complete discussion of various ways to optimize storage for HVD.

GENERAL GUIDELINES TO MINIMIZE COST PER DESKTOP

In addition to performing proper sizing, follow these best practices to reduce cost:

- Use SAS or FC disks for virtual desktops. An HVD environment is typically performance constrained. An analysis shows that cost per IOP is significantly less expensive on SAS or FC than on SATA, even after optimization technologies such as Flash Cache and WAFL are applied.
- Separate user home directory data from the virtual desktop data and then put the user data on SATA. User data tends to be capacity intensive but requires very little performance. Separating user data from the virtual desktop data has the added benefit of simplifying patching, backup, and recovery.
- Align the file system and storage layer. Misalignment incurs significant additional reads and writes. The following section provides a more detailed discussion of misalignment.
- Deploy a fresh virtual desktop image because it usually requires less storage than converting from a physical desktop (that is, P2V).
- Optimize the desktop for HVD. Follow the virtualization vendors' best practices for guest optimization in the following guides: [VMware View Optimization Guide for Windows 7](#), [Windows XP Deployment Guide](#), [Citrix Windows 7 Optimization Guide](#).
- Disable all defragmentation, indexing, and limiting search tools to minimize the disk I/O from the VM.
- Disable or stagger running tools such as antivirus and Windows updates that perform simultaneous downloads to prevent massive I/O spikes back to the disk that require more storage to support the VDI images.

COMMON CULPRIT OF STORAGE OVERCONSUMPTION: MISALIGNMENT

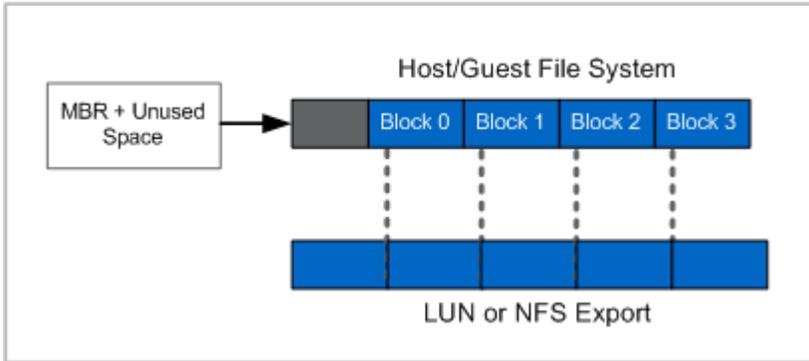
In any virtual environment, there are a number of layers of abstraction between physical disks and the VM's virtual disk. Each layer in turn is organized into blocks to make the most efficient use of storage. The focus is not on the size of the block, but rather on the starting offset. Failure to align the file systems results in a significant increase in storage array I/O in order to meet the I/O requirements of the hosted VMs. Misalignment results in every I/O operation executed within the VM to require multiple I/O operations on the storage array. Customers might notice this impact in VMs running high-performance applications or with less than impressive storage savings with deduplication or excessive CPU load on the storage array. Figure 8 illustrates file-level misalignment.

To avoid unnecessary load caused by extra reads and writes, the starting offset of a file system should line up with the start of the block at the next layer down. This guideline is not unique to NetApp; it applies to any storage vendor. It is a simple byproduct of legacy partitioning schemes. For the full explanation of

disk alignment in virtual environments, see [TR-3747: Best Practices for File System Alignment in Virtual Environments](#).

Traditionally, disk utilities such as GNU fdisk are used to alter the real geometry of a disk drive to partition the disk in a more efficient manner. Unfortunately, this does not usually bode well, because the default values used by tools such as fdisk do not align the offsets properly for use in a virtual environment.

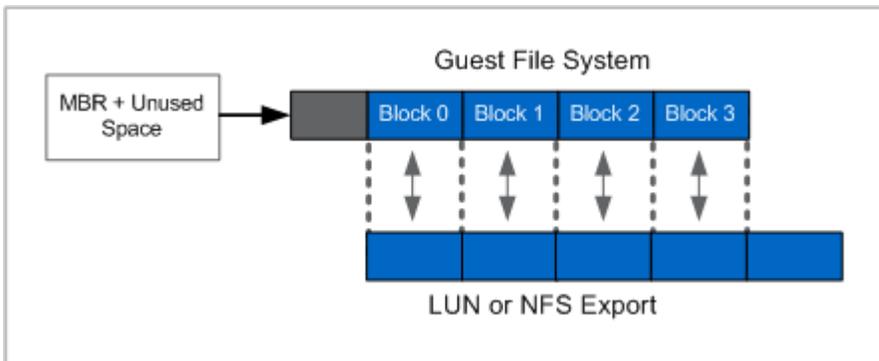
Figure 8) File-level misalignment.



NetApp uses 4KB blocks ($4 \times 1,024 = 4,096$ bytes) as its basic storage building block size. Write operations can consume no less than a single 4KB block and can consume many 4KB blocks depending on the size of the write operation. Ideally, the guest/child operating system (OS) should align its file system(s) such that writes are aligned to the storage device's logical blocks. The problem of unaligned LUN I/O occurs when the partitioning scheme used by the host OS does not match the block boundaries on the back-end storage array.

Figure 9 shows the proper alignment of the guest file system, through the host node file system, and down to the blocks on the NetApp FAS controller. Although a LUN is used in the environment, file system alignment is required regardless of the storage protocol used (NFS, iSCSI, FC, or FCoE).

Figure 9) Proper alignment of guest file system.



VMs running a clean installation of Microsoft Windows 2008, Windows 7, or Windows Vista operating systems automatically have their starting partitions set to 1048576. By default, this value does not require any adjustments.

If the customer is not using Microsoft Windows 2008, Windows 7, or Windows Vista, the alignment process should not be considered optional. Misalignment at a high level results in decreased utilization. Make sure you align your storage partition for your golden images before you install the operating system. If the misaligned golden images have already been created, you can run the NetApp MBRAAlign tool to properly realign it.

For more information regarding file system alignment, see [TR-3749: NetApp and VMware vSphere Storage Best Practices](#) and [TR-3747: Best Practices for File System Alignment in Virtual Environments](#). Realignment is a fairly intensive process and should be done only as a last resort. NetApp strongly recommends that you avoid misalignment altogether in the design process if possible.

In certain situations, file system misalignment cannot be avoided. In these cases, an adjustment must be made to the desktop profile, and extra overhead IOPS must be added and accounted for in sizing. The effects due to misalignment are difficult to predict exactly, which complicates the profiling exercise. Misaligned I/O can generate an increased amount of disk reads due to a VM write operation as well as increasing the CPU load on the storage controller. Therefore, NetApp recommends sizing storage systems that host misaligned desktops with extra CPU overhead and an increased number of disk drives to support the misalignment. Your misalignment overhead might vary depending on how risk averse you want the desktop profile to be. For a more detailed discussion of sizing for misalignment, consult the NetApp sizer and your NetApp SE or PS.

COMMON CULPRIT OF STORAGE OVERCONSUMPTION: BLACKLIST APPLICATIONS

When profiling storage for virtual desktops, it is important to learn as much as possible about the applications contained on the desktop as well as the management practices of the desktop team. Certain activities that are trivial on physical desktops can have a large effect on virtual desktops. For example, consider an antivirus application downloading a 15MB update package. Most users would not notice their physical desktop downloading the package, because they each have their own dedicated hard drive. Therefore, each 15MB is streamed to one disk spindle. Consider an HVD deployment on NetApp storage where storage has been centralized and consolidated. A 15MB download distributed across 1,000 desktops suddenly becomes 15GB of disk writes targeted at a set of centralized spindles. This could adversely affect the storage performance and the end-user experience.

Another area of concern is any application that tries to optimize disk performance on the desktop. Most of these applications were written with the physical desktop in mind, and they often attempt to arrange the data on the disk with the assumption that a single disk platter is in use. With HVD, however, the desktop image is distributed across a large number of disk platters and interleaved with other disk images. In this case, having the application move or “optimize” disk layout is unnecessary and simply generates extra I/O on the disk subsystem.

Yet another area of concern is applications that search or index the contents of the desktop. Many applications today attempt to build an index of the desktop contents as a background operation. This enables faster searches for the end users when they search for content. Again, this is possible in a physical environment where the hard drive is underused. In an HVD environment, however, this indexing can be a major source of storage I/O affecting performance.

To summarize, be aware of the following applications and size for them accordingly:

- Autoupdaters (antivirus, application update)
- Optimizers (disk defragmentation, disk optimization)
- Indexers (Microsoft Office index, Google toolbar, desktop search)

Use the NetApp custom report in Liquidware Lab’s Stratusphere software to identify these “blacklist” applications. After these applications are identified, the next step is to implement new ways of delivering these applications without incurring the performance cost. You might want to use application streaming to deliver some applications. Some security vendors, such as Trend Micro, have been making innovations to their antivirus scanning to minimize the impact of scanning and definition update. In some cases, you might choose not to move these “blacklist” applications to the HVD environment at all.

4 ACKNOWLEDGEMENTS

The following people contributed to the creation and design of this guide:

- Lee Dorrier, Director, NetApp
- Chris Gebhardt, Technical Marketing Engineer, NetApp
- Neil Glick, Technical Marketing Engineer, NetApp
- Dnyanesh Khare, Product Manager, Liquidware Labs
- Mike Shea, Global Solution Lead, NetApp

5 FEEDBACK

Send an e-mail to xdl-vgibutmevmtr@netapp.com with questions or comments concerning this document.

6 VERSION HISTORY

Version	Date	Document Version History
Version 1.0	March 2011	Original document

7 REFERENCES

- [Citrix Windows 7 Optimization Guide](#)
- [TR-3747: Best Practices for File System Alignment in Virtual Environments](#)
- [TR-3749: NetApp and VMware vSphere Storage Best Practices](#)
- [VMware View Optimization Guide for Windows 7](#)
- [Windows XP Deployment Guide](#)

NetApp provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.