



Education

Storage Optimization for VDI

Russ Fellows
Evaluator Group, Inc.

- ◆ The material contained in this tutorial is copyrighted by the SNIA unless otherwise noted.
- ◆ Member companies and individual members may use this material in presentations and literature under the following conditions:
 - ◆ Any slide or slides used must be reproduced in their entirety without modification
 - ◆ The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- ◆ This presentation is a project of the SNIA Education Committee.
- ◆ Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- ◆ The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

➤ Speaker

- ◆ Russ Fellows , Senior Partner, Evaluator Group

➤ Description

- ◆ We will explore the critical storage features needed for Virtual Desktop Initiatives (VDI). The focus will be on storage system features required to meet the performance, and the critical TCO levels necessary for a successful VDI project. The recommendations are taken from hands on experience implanting VDI projects in mid-sized and large IT environments. Particular attention will be placed on understanding the storage implications of different VDI implementations and how this impact performance and overall project cost. The emphasis will be on practical configuration options and how best to leverage specific storage features in order to maximize storage performance and efficiency.

- Implementing VDI
 - ◆ Benefits
 - ◆ Real World Issues with VDI
- Actual VDI Storage Workload
- Storage Considerations
 - ◆ VDI Design
 - ◆ Storage Features
- Recommendations
- Storage Performance for VDI
- More Resources

IMPLEMENTING VDI

➤ May run on server or locally

- › Server based is typical – uses centralized IT resources
- › Local mode is a new – runs using local hypervisor

➤ Different modes for VDI images

- › Persistent – maintain changes (similar to existing use)
- › Pooled – are not modifiable (OS reverts to original state)

➤ Pros and Cons

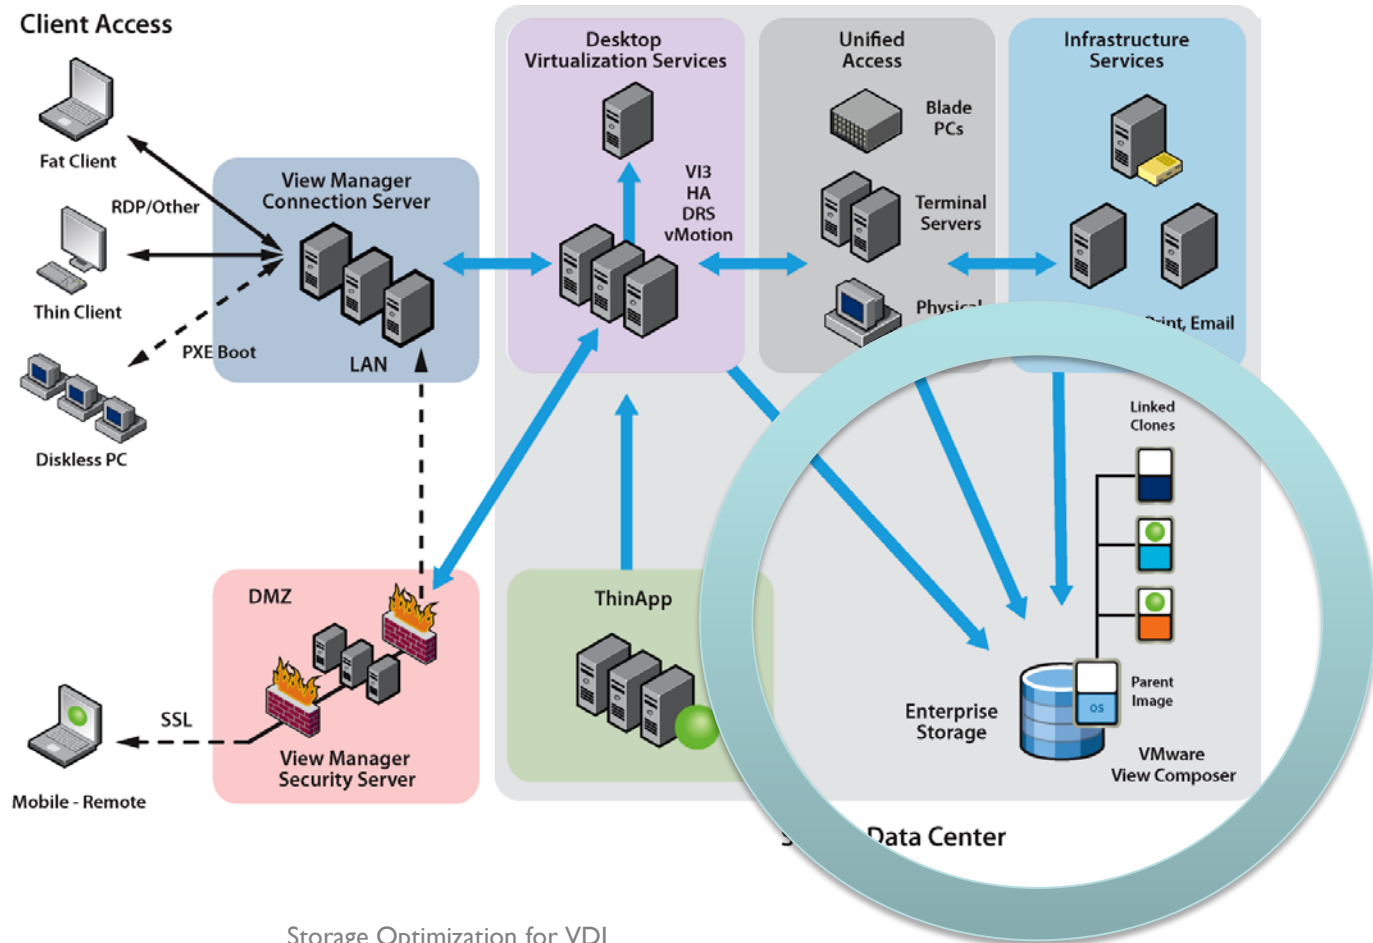
◆ Advantages:

- › Centralized administration (Reduced Opex)
- › Improved Security Controls (access, data restrictions)

◆ Disadvantages

- › May cost more than traditional desktops, storage often an issue

Storage is the Heart of VDI

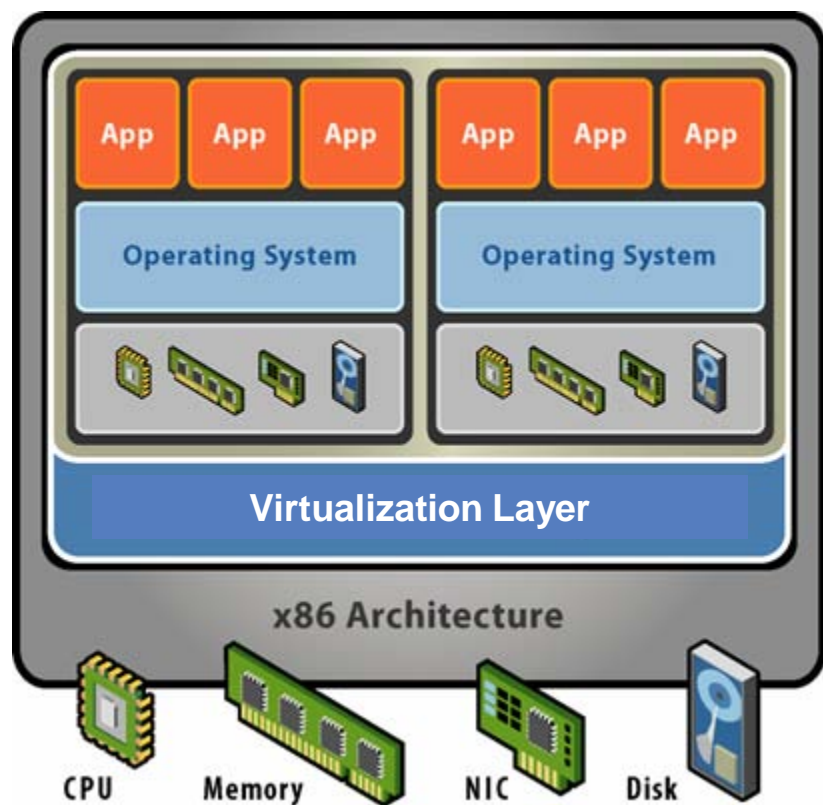


Barriers to VDI Adoption

- **May require network connection**
 - ◆ Local mode removes this limitation
- **Capital Investment (particularly storage)**
 - ◆ Storage costs may exceed expectations
 - ◆ Capital costs typically exceed traditional approach
 - ◆ Storage configuration is critical to cost and performance
- **Concerns over technology or approach**
 - ◆ Users dislike lack of control (particularly with pooled)
 - ◆ Concerns over performance and flexibility

The impact of VDI on Storage

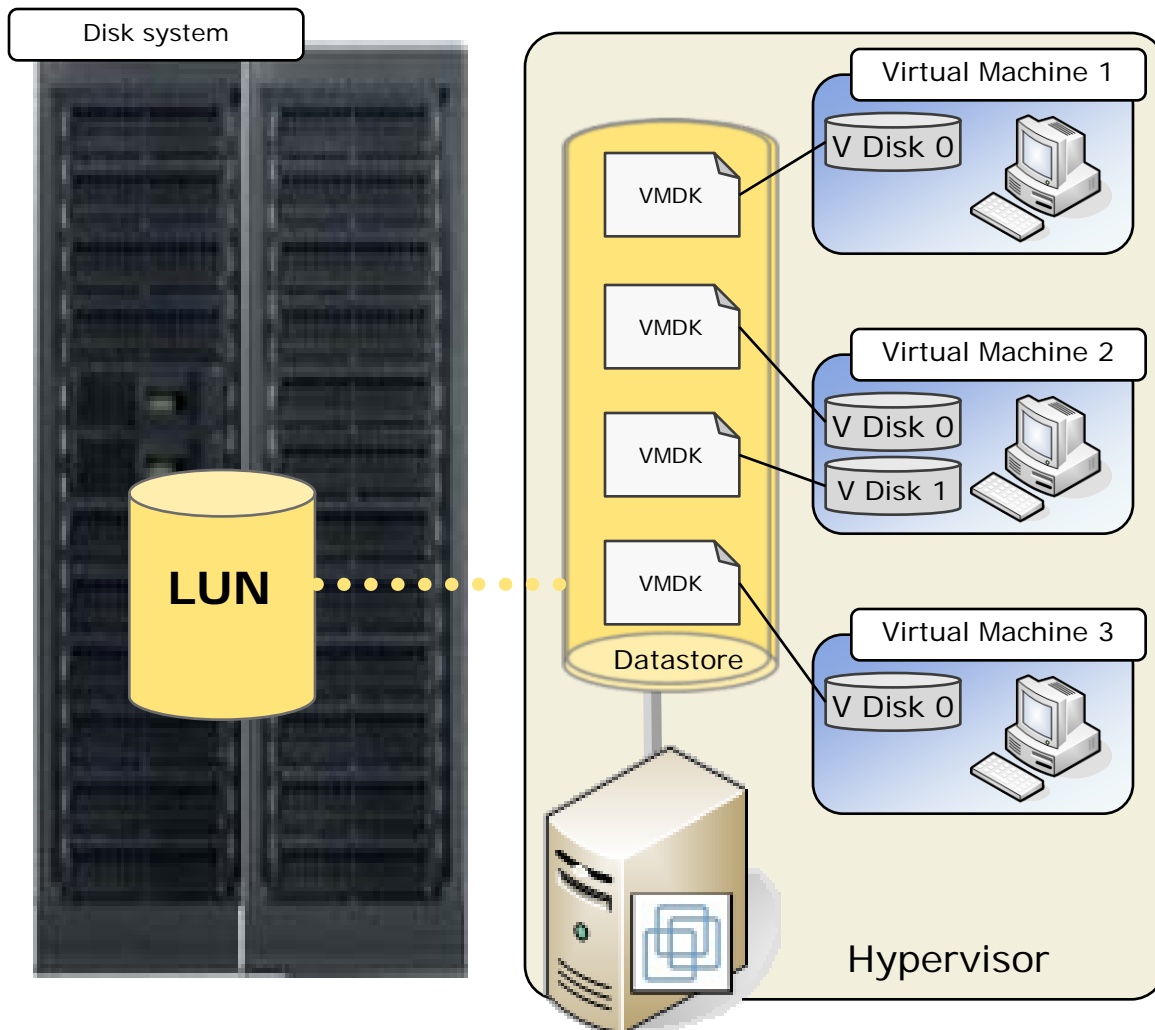
VDI STORAGE CONSIDERATIONS



Virtualization Technology

- **Separation** of OS and hardware
- **Encapsulation** of OS and application into VMs
- **Isolation**
- **Hardware independence**
- **Flexibility**

OS Using a Hypervisor



Components

Bare Metal Hypervisor
- Physical server platform

Shared Storage
- FCP, iSCSI, or NFS

Datastore
- Shared storage configured to store VMware data

Virtual Machine (VM)
- Logical collection of hardware resources such as CPU, memory, network, and storage

Virtual Disk (VMDK File)
- Virtual SCSI attached storage for a VM

VDI Workload Considerations

➤ Virtual Server Performance

- ◆ Random I/O performance
- ◆ Wide Striping
- ◆ Auto Tiering
- ◆ SSD / Flash Support

➤ VDI Performance

- ◆ Writeable clones
- ◆ SSDs as storage (not cache)
- ◆ Lock master image on SSD
- ◆ Cache (only somewhat useful)


➤ Nice to Have

- ◆ 1,000's of Snapshots
- ◆ Data Deduplication (limited use for VDI)
- ◆ Thin Provisioning
- ◆ Automated Provisioning (especially for VDI)


➤ Not Useful

- ◆ Tiering without clones

Duplicate Data:
OS and App
> 90% same



All data -
Thru
Server



Storage - each App and
OS contending for resources

- Same problem as virtual machines
- Except – multiplied
 - ◆ Architecture is critical
 - ◆ More users
 - ◆ Likelihood of “boot storms”
 - ◆ Simultaneous application startup
 - ◆ Protection and recovery – frequency and user-level control
 - ◆ Provisioning x 1000

➤ Desktop Performance

- ◆ Must be similar to existing, or risk project failure
- ◆ Requires CPU, Memory, Network and Storage
- ◆ Performance during steady state, and I/O Storms

➤ I/O Storms

- ◆ Boot Storms
- ◆ Login Storms
- ◆ Virus Scanning
- ◆ Patching / re-imaging / refreshing image

➤ Cost

- ◆ Requires careful architecture to achieve low \$ / VDI user

➤ Architecture for I/O Optimization

- ◆ Persistent vs. Non-Persistent : Understand Storage Impact
- ◆ When to use Hypervisor tools vs. Storage tools

➤ Requires use of Solid State Storage

- ◆ On a \$ / iops basis, SSD and Flash less expensive
- ◆ Use of Solid State limited to master image
- ◆ Use Traditional storage for changes (i.e. writes)

➤ Caching Appliances may Improve Performance

- ◆ Virtualization and VDI specific Storage
 - › Vendors include: (Whiptail, I/O Turbine, Virsto, Atlantis Computing, Nimble Storage, Tintri, etc.)

➤ Storage System has a Significant Impact

- ◆ Large cache (both read and write)
- ◆ SSD for golden image
- ◆ Writeable clones
- ◆ Wide Stripping
- ◆ Nice to Have:
 - › Many snapshots (for data protection)
- ◆ Little Impact:
 - › Deduplication, (if using pools, or storage clones)

Real World VDI Storage Workload Characteristics

VDI STORAGE WORKLOAD

Storage for VDI

➤ Limited Guidance for Storage

- ◆ Capacity (varies based on pooled / persistent)
 - › MS recommends 20 GB / Win 7
- ◆ Performance (particularly IOPS)
 - › Use 8 – 12 iops as an average

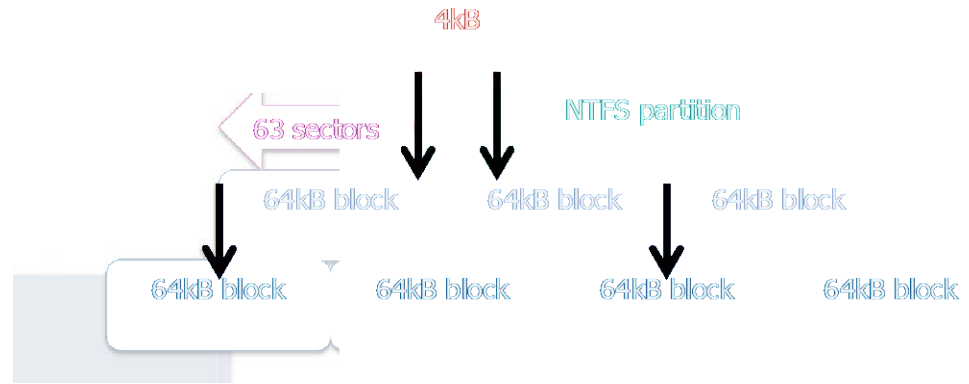
➤ Issues:

- ◆ Workload Claims are Inaccurate
 - › Claims of 8 IOPS / client are highly generalized
 - › I/O rate varies, heavy-user profiles are closer to 12 iops
 - › Read / write ratio's shift towards read for heavier use
- ◆ A real VDI I/O is different than a synthetic I/O
 - › A real I/O is often 2 MB, synthetic i/o tests use 512 Bytes
 - › I/O's are not small blocks, or a single block size
- ◆ “Law of Averages” doesn't work
 - › Queuing theory dictates 3X headroom needed to manage peaks

OS Image Layout Issues


➤ Default Windows Alignment

- ◆ Starts at 63 Sectors
- ◆ Causes all subsequent blocks to be unaligned



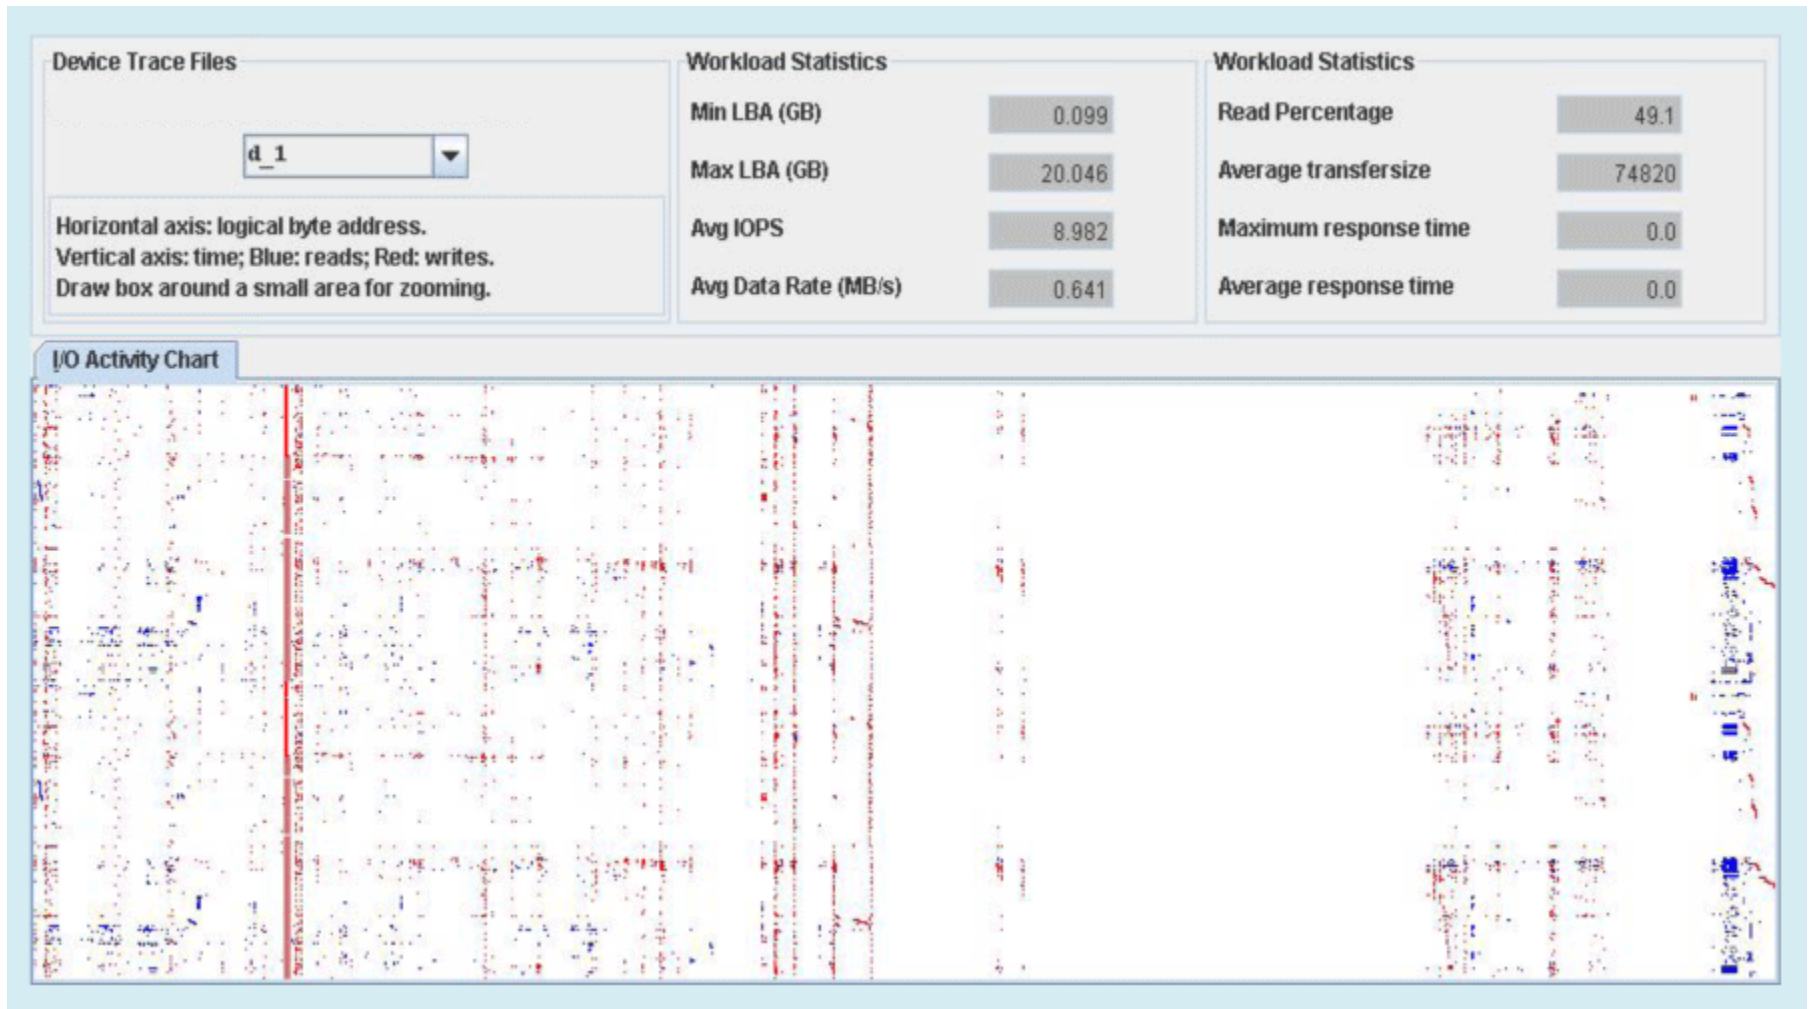
➤ Recommended Layout

- ◆ Start at sector 128
- ◆ All subsequent I/O's should be aligned




- A “Standard” VDI workload has following characteristics:
 - ◆ Range from 5 iops – 20 iops
 - › Depends upon jobs running, etc.
 - › MS Office apps generate relatively little I/O after launch
 - › 7z, Media Player, Web browser and Outlook create greatest I/O (in order)
 - ◆ Read / Write mix is about 50 – 50
 - › Lighter workloads are skewed towards more writes
 - ◆ 4K block transfers are the most common for writes
 - ◆ 32K block transfers are the most common for reads
 - ◆ 1M block transfers move the most data for writes
 - ◆ Reads move the most using equal # of (2M, 1M, 32K and 64K)
 - ◆ Constant writes at about 5 GB sector for all 20G partitions
 - ◆ Heavy reads around 17 – 20 GB sectors
- Graphs depicting information on following 3 slides

I/O by Address (x) and Time (y)



Distribution of I/O Transfer Size

I/O Distribution by Transfer Size
(Controllers: (Real Controllers), All devices, Detail) (avg: 88447 bytes)

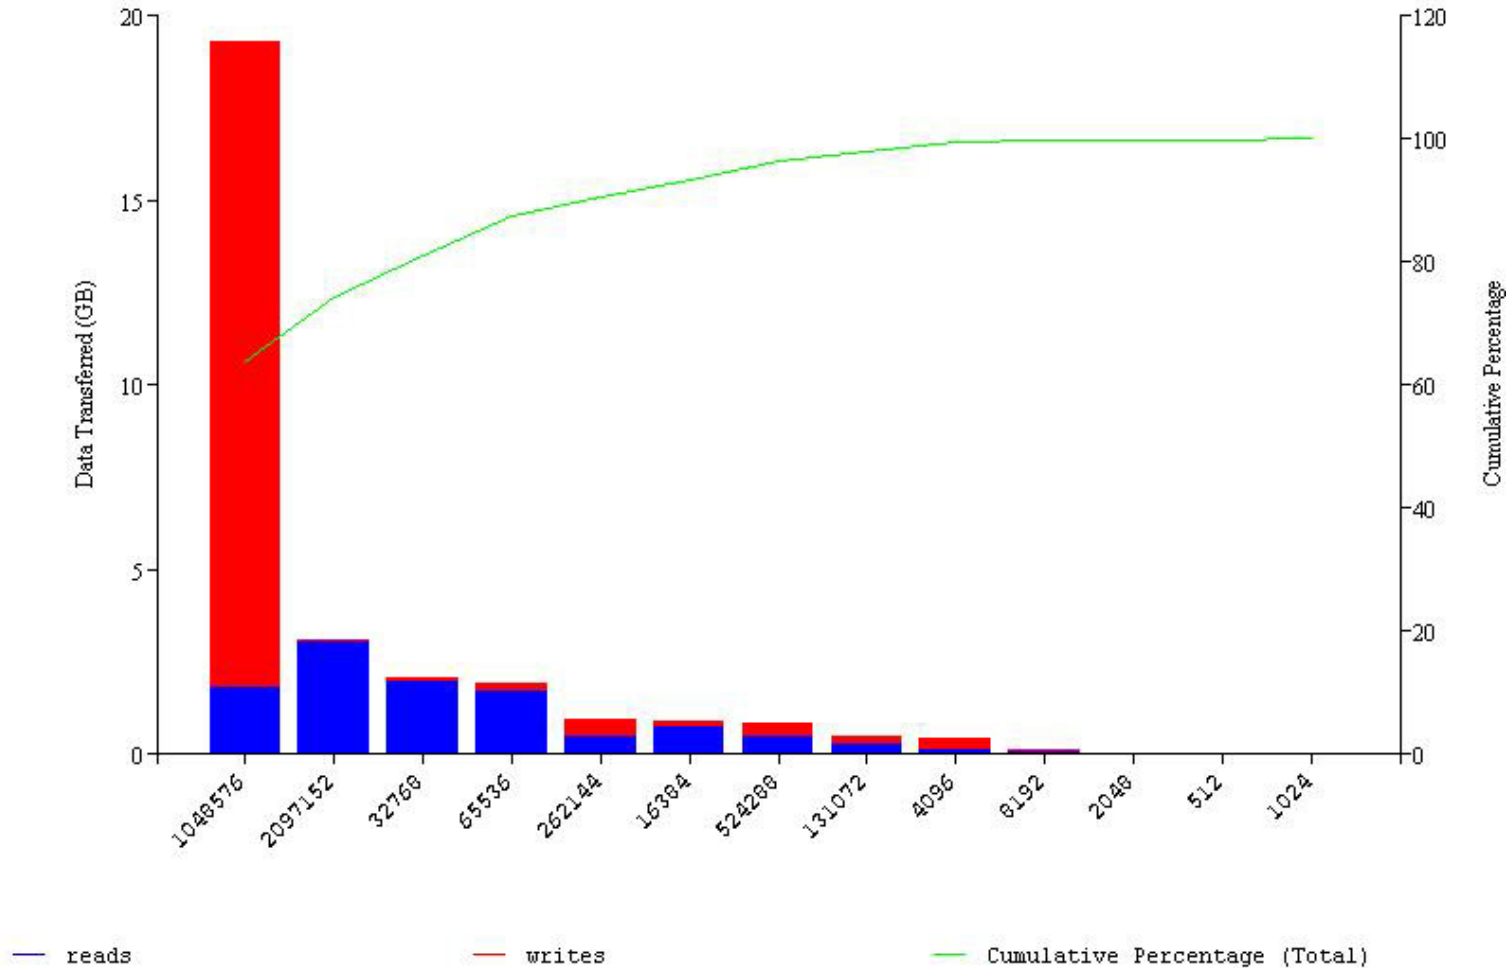


Storage Optimization for VDI

© 2011 Storage Networking Industry Association. All Rights Reserved.

Distribution of Data Transferred

I/O Distribution by Data Transferred
(Controllers: (Real Controllers), All devices, Detail) (avg: 88447 bytes)



STORAGE PERFORMANCE FOR VDI

- Existing deployments show storage is the critical item for performance
- Difficult to rationalize vendor claims
 - ◆ Vendors claim to support a “VDI user” for a specific \$ value, but there is no common metric for configuration or workload used to make those claims in a VDI environment
- Existing VDI tools test entire system
 - ◆ Requires extensive server, network and environment setup
 - ◆ Considerable expense (100’s of servers)
- Existing storage benchmarks do not recreate VDI workload

➤ The Issues:

- ◆ Proof of Concept can be costly
- ◆ Typically requires all server, network and storage resources
- ◆ Extensive setup time (can be multiple weeks)
- ◆ Most performance tools require full environment
 - › Login VSI, View Planner, VMware RAWC (only through partner)

➤ Recommendation:

- ◆ Storage VDI Benchmark
 - › VDI-IOmark

VDI Storage Performance Testing

	VDI-IOMark	Login VSI	RAWC	SPC
What it tests	Storage	Entire System	Entire System	Storage
Workload	100% real VDI	100% real VDI	100% real VDI	Non VDI
Cost	No cost for Users – Vendors pay	Cost to license	VMware partners only	Cost to license & publish
Equipment Required	Low	High	High	Low
Setup Time	Low	High	High	Low

Benefits of VDI Sizing Tools

- Accurate Results (Login, RAWC and VDI-IOmark)
- Reduce setup and equipment (SPC and VDI-IOmark)
 - ◆ Reduced Time – 1 day setup vs. 1 week
 - ◆ Reduced Expenses – 30x less server equipment
- Example: 5,000 seat VDI POC (VDI-IOmark)
 - ◆ Traditional POC (actual, Login VSI, RAWC, etc.)
 - › 32 Users per server (appx. 160 servers)
 - 160 Servers @ \$10K (96 GB, 12 cores, multiple I/O ports)
 - › Cost = \$1.6 M for servers
 - ◆ VDI-IOmark 5,000 user test
 - › 1,000 Users per server (5 servers)
 - 5 Servers @ \$10K (96 GB, 12 cores, multiple I/O ports)
 - › Cost = \$50 K for servers

A Storage Benchmark for VDI

- A storage specific benchmark for VDI
 - ◆ Tests storage only
- Storage Agnostic
 - ◆ Supports any storage supported by hypervisor
- Based on actual VDI workloads
 - ◆ Uses “I/O replay” to simulate storage I/O patterns
 - ◆ Benchmark drivers use real-world workloads
- Far less resources (an order of magnitude lower)
 - ◆ Lower capital requirements by 10x
 - ◆ Reduced test time 10x less setup
 - ◆ Each server can test up to 1,000 VDI users
 - › 12 CPU, 96 GB RAM, with multiple I/O ports

➤ Benchmark creation

- ◆ Real world VDI environment used for data capture
- ◆ Actual I/O captured in VDI configuration and workloads (RAWC based)

➤ Benchmark run

- ◆ Utilize driver to replay workloads (I/O replay)
- ◆ Benchmark driver (the test harness) does not require applications
- ◆ Benchmark driver runs under multiple guest OS
 - › Linux supported initially, Windows TBD

➤ Result reporting

- ◆ Results indicate number of users supported
- ◆ Benchmark runs and results audited by EGI for consistency
- ◆ Storage configuration options are included in report

RECOMMENDATIONS

- Use clones when possible
 - ◆ Either pools, or linked clones on VM
 - ◆ Or, utilize storage writeable clones
- Separate Master Image from Other Data
 - ◆ Place Master Image on Solid-State
 - › Ensure master image resides on Solid-State
 - › Either VM or Storage Cloning may work
 - ◆ Maintain differences in separate disk area
 - › Swap / page area, User profile data, etc.
 - › This data should be places on spinning media

- De-dupe does not offer practical nor operational benefits when using Floating or Persistent Pools with Linked Clones.
- De-dupe may be beneficial in environments that use Persistent Full Clones 100%, however Full Clones and Manual Pools are not the way to move forward with your VDI strategy, specially in DR situations.
- \$ per IO is focus, \$ per GB irrelevant
- VMware Linked Clones and Citrix MCS provide many performance and control advantages

- Separate OS and User data (good advice generally)
- Utilize some form of writeable clones
 - ◆ VMware Linked Clones or Citrix
 - ◆ Storage writeable clones
 - ◆ Utilize storage utilities if available to improve performance
- **MUST** utilize flash for some portion of data
 - ◆ Using clones concentrates I/O's for reads
 - ◆ Use Solid State for golden image read data (10 – 20 GB)

- Please send any questions or comments on this presentation to SNIA: tracktutorials@snia.org

**Many thanks to the following individuals
for their contributions to this tutorial.**

- SNIA Education Committee

Joseph White

THE END

Additional Resources

APPENDIX

➤ Capital Costs

- ◆ Storage capital costs
- ◆ VDI requires enterprise class storage
 - › Replace 1,000 inexpensive desktop drives with 500 enterprise drives in HA RAID with storage network

➤ Management Issues

- ◆ Creating clones can be time consuming
- ◆ VMware clones take significant resources to create

- The holy grail of the master golden image
 - ◆ Install all apps, recommended patches
 - ◆ Perform all recommended changes (250+)
 - ◆ Shrink partition (releases allocated free space)
 - ◆ Create Snapshot
 - ◆ Do not to use Thin Provisioned disks for Parent VMs, instead create them eagerzerothick.

➤ Persistent vs. Non-persistent

- ◆ As the name implies, persistent saves changes
- ◆ Both XD and View have the capability of running persistent (private) desktops or standard (pooled) desktops
- ◆ Persistent images unique to a specific end user and they may make updates and install applications.
- ◆ With Pooled (standard) each instance runs from a single OS image that is read only. Changes are stored for the session only and the image reverts to the golden image upon a user logoff/logon.
- ◆ VMware supports Linked clones with persistent images
- ◆ Xen MCS supports a base image for each catalog config

- VDI instances should use either
 - ◆ Pooled / shared instances
 - ◆ Persistent / private images with clones
- Persistent Desktops should be an exception to address specific user or application needs
- Data de-duplication has limited benefit in these cases
- Thin provisioning has limited benefit also
- VDI storage costs are related to IO operations, not capacity

Example: Linked Clone

