



Storage Performance and Desktop Virtualization

James Candelaria, CTO



Company/Presenter Snapshot

- WHIPTAIL incorporated in (2009) direct response to the IO challenges faced by enterprises in deploying desktop and server virtualization technologies.
- WHIPTAIL's technology stack drives down the cost and complexity of high performance storage by offering enterprise class family or arrays that dramatically increases end user productivity and slashing the cost per IOP.





Over 4 years of VDI successes



XenDesktop

XenDesktop

1,000 XenDesktop Users



2,500 VDI Users



VMware View



2,000 XenDesktop Users



1,500 VDI Users



3,000 VDI Users



The Pensions Trust



8,000 XenDesktop Users



XenDesktop



MOTOROLA

1,000 XenDesktop Users





1,000 XenDesktop Users



Santa Clara, CA
August 2012

Unaccounted for Performance

- CPU & Memory is easy to measure 
- Disk resources are not 
- SBC solutions only had one OS arbitrating disk access, while VDI has THOUSANDS of INDEPENDENT operating systems with:
 - Unshared read caches
 - Virtual memory footprint for EVERY user
 - Easy user runs their own AV, security suite, inventory agent, etc.
- New class of problems (storms) introduced:
 - Boot storms
 - Patch storms


Solving the problem(s) – Host side – PCIe/Flash drive DAS

- PCIe
 - Very fast and low latency
 - Expensive per GB
 - No redundancy
 - CPU/Memory stolen from host
- Flash SATA/SAS
 - More cost effective
 - Can't get more than 2 drives per blade
 - Unmanaged can have perf/endurance issues




Solving the problem(s) – Array based cache/tiering

- Array flash cache
 - Typically read only
 - PVS already caches most reads
 - Effectiveness limited by storage array designed for hard disks
- Automated storage tiering
 - “Promotes” hot blocks into flash tier
 - Only effective for READ
 - Cache misses still result in “media” reads





Solving the problem(s) – Flash in the traditional array

- Flash in a traditional array
 - Typically uses SLC or eMLC media
 - High cost per GB
 - Array is not designed for flash media
 - Unmanaged will result in poor random write performance
 - Unmanaged will result in poor endurance

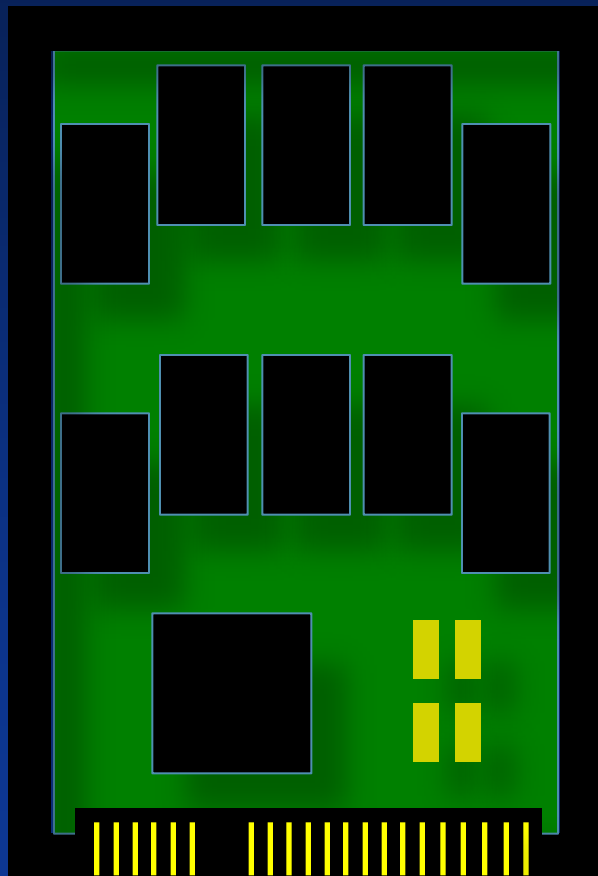


Solid State Fundamentals: *HDD Write Process Review*




A physical HDD is a bit-addressable medium! Virtually limitless write and rewrite capabilities.

Solid State Fundamentals: *Flash Writes Process*



2MB ERASE BLOCK



4K data
blocks

1. Erase block contents are read to a buffer.
2. Erase block is wiped (aka “flushed”).
3. Buffer is written back with previous data and any changed or new blocks – including zeroes.


- Why does this matter?
 - Each cell has physical limits (dielectric breakdown)
 - Time to erase a block is non-deterministic (2-6ms)
 - Program time is fairly static based on geometry
 - Failure to control write amplification *will* cause wear out in a short amount of time
 - Desktop workload is one of the worst for write amplification
 - Most writes are 4-8KB



Understanding Random Write Performance

- Write amplification not only causes wear out issues, it also creates unnecessary delays in small random write workloads
- What is the point of higher cost flash storage with latency between 2-5ms?

Managing the write process



Write requests are flushed the
RAID layer as full stripes and
to media as full ERASE
BLOCKS




Deployment Scenarios

- Pod based:
 - 2500-5000 users per pod
 - Dedicated blade chassis per pod
 - N+1 pods
 - 1-2 PVS servers per pod
 - Write cache and hold image storage on WHIPTAIL
 - Profiles/home directories everywhere
- Centralized
 - X number of blade centers to support users
 - 1 INVICTA chassis per deployment
 - Gold images mirrored inside of INVICTA
 - 2-4 PVS servers per deployment
 - Write cache and gold image storage on WHIPTAIL
 - Profiles/home directories everywhere



Customer Example

IVENT VDI Implementation using Flash



Defense material
Organisation
Ministry of Defense

Operations/Infrastructure
and Services
SBR
Cor van 't Hoff
Sr. Designer



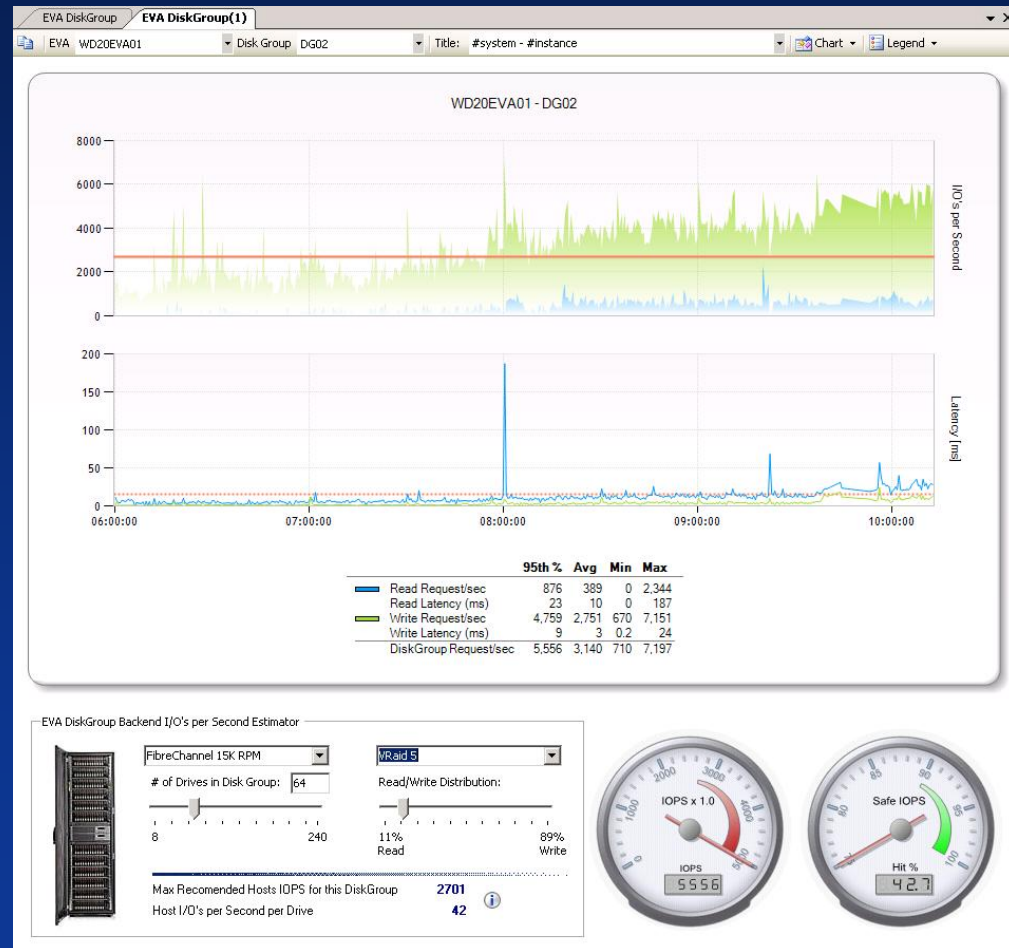
Defense Material Organisation IVENT

- Our key areas
 - Support for over 55.000 work stations
 - Archives of more than 13 kilometers in length
 - Research & innovation Center
 - Nation-wide fibreglass network at out disposal
 - Two command centers with 24/7 security
 - Three secured data centers
 - Owner of independent back-up center
 - Defense Service Desk en Operations Rooms
- Our certifications
 - Oracle Certified Partner
 - Certified as SAP Competence Centre
 - Microsoft Gold Partner
 - ISO certification for the IVENT production and services process



Problems using traditional storage for VDI

- Running VDI on our existing hard disk array didn't take off
- During testing less than half of the projected 2600 VDI users could log in
- IOPS were more than the array could handle
- Latency would rise to unacceptable levels

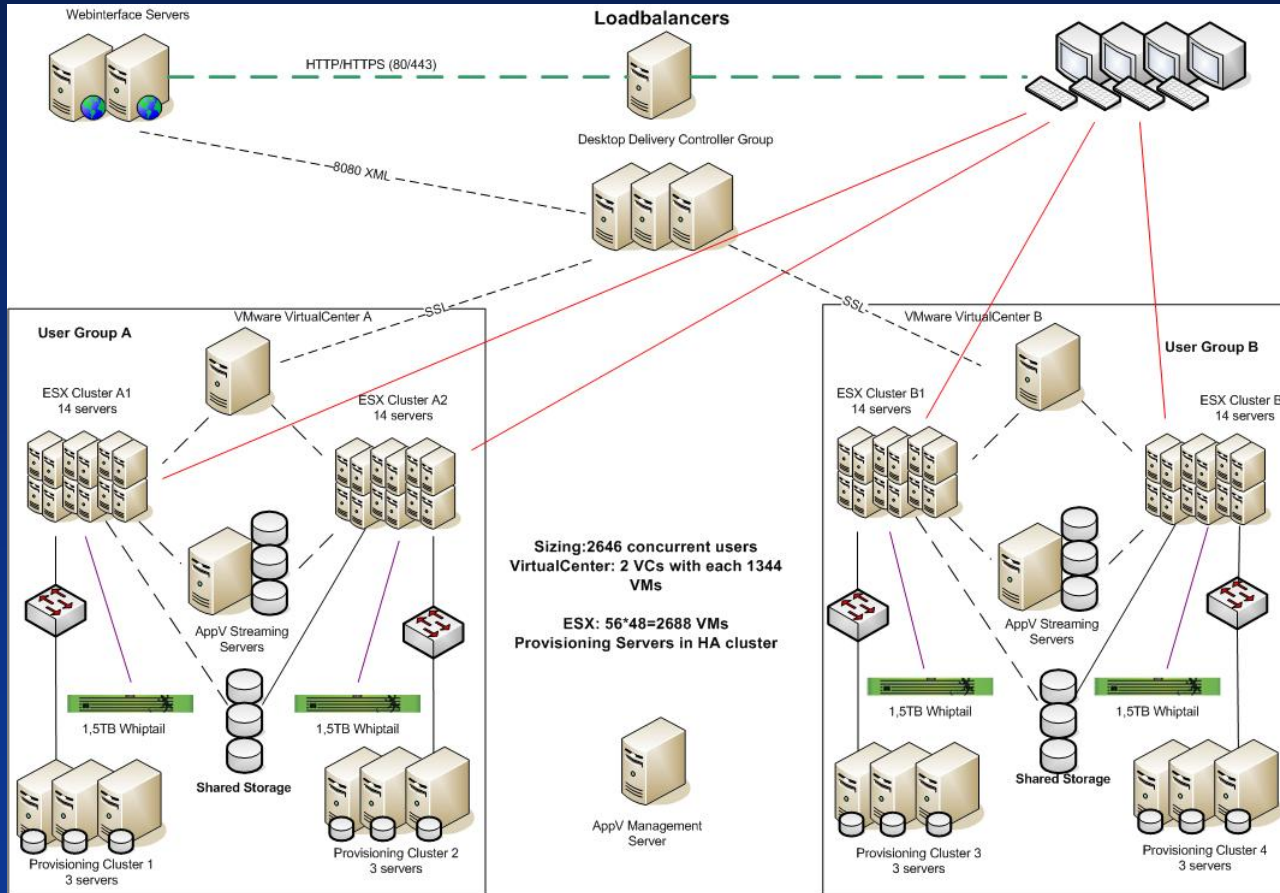




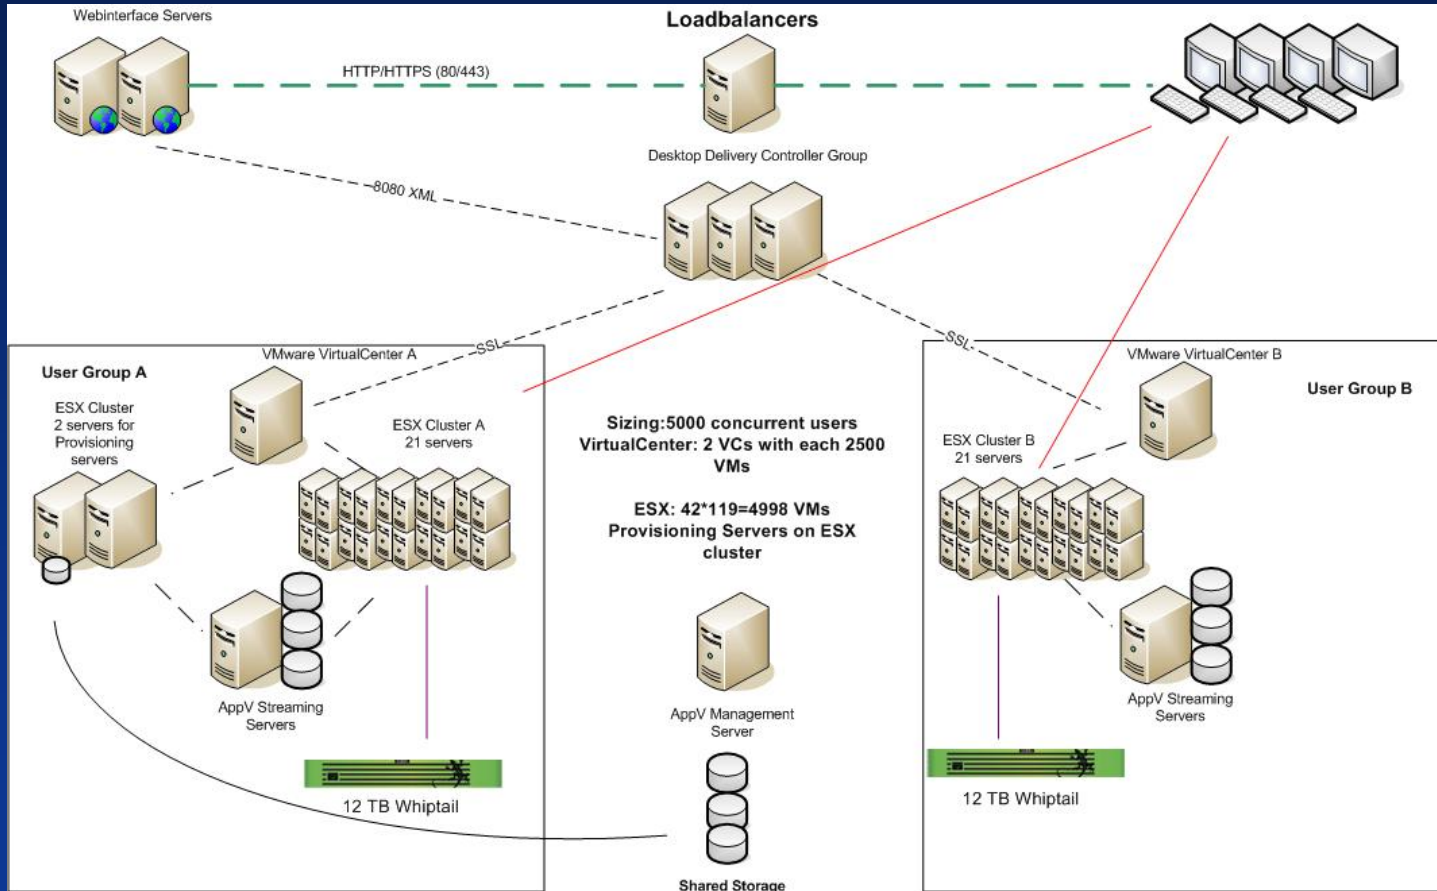
Flash Memory Solutions

- Add extra disks to the array
 - Array was already full
- Buy other hard disk arrays
 - Expensive and would require lots of power and space
- Use flash based PCIe cards
 - No shared storage
- Purchase a flash based Whiptail
 - ✓ Requires less power and space
 - ✓ Shared storage
 - ✓ No I/O bottleneck for the foreseeable future

2650 concurrent user environment



MoD 5000 concurrent user POD





Current and planned scale

- Current
 - 1 environment 2650 concurrent users
 - 1 environment 5000 concurrent users

- Planned this year
 - 1 environment 2650 concurrent users
 - Scale the 5000 concurrent user POD to 7000 users
 - Deploy 2 PODs



Thank you

