

FILE IO BOTTLENECK

PRESENTED BY:

JAMES CANDELARIA

CCEA, VCP, MCSE, CCNA

CTO


WHIPTAIL TECH

AGENDA



SO WHAT IS VDI?

- Many different definitions
- Hosted virtual desktop the most common
 - Client OS (Win7/XP) hosted on datacenter hypervisor (ESX/Xen) + MGMT layer
- Terminal services sessions also common
 - Multi-User Windows Server OS + MGMT layer



WHY VDI?

VDI DESIGN CONSTRAINTS

- CPU
 - Good tools exist to profile and plan for, typical consolidation ratio of 10+ users per core
- Memory
 - OS tools provide consumption data, 1-2GB of RAM per user (depending on OS)
- Storage performance
 - Highly variable, hypervisor and deployment model dependent

IO PROFILE (STEADY-STATE)

- Vastly different than typical datacenter profile
 - Datacenter – 80% read 20% write
 - VDI – 50/50 and very small request size
- Incredibly random
 - Even sequential IO looks random to the array after being chopped up by the hypervisor
- Massive IO volume
 - Terabytes per day of IO for even a small < 5,000 seat deployment

IO PROFILE (BOOT/LOGIN STORM)

- Massively random workload can crush traditional arrays with over 1GB/s of random IO.
- Environments often must be rebooted after
 - OS Patching
 - Infrastructure maintenance
 - Hypervisor maintenance operations
- Login storms
 - Typically at least twice per day
 - Morning
 - After lunch

MAGNETIC STORAGE ISSUES

- 2,000 users requires 40,000 IOPs (minimum)
 - Or 200 15K drives
 - Drawing 4.5 KW
 - 51 Rack units of space
 - > \$500,000 or \$250 per user
- Caches are largely ineffective
 - Write caches fill quickly with this many users
 - Read caches not large enough for 2000 users

HOW DOES FLASH HELP?

- Pure flash array can supply GB/s of random performance
 - Boot / login storm mitigation
- Well designed flash array can deliver sustained read/write performance in excess of 100,000 random IOPs
 - Steady state

WHAT TO WATCH FOR?

- Most writes are small and random
 - Small random writes can be the enemy of flash
 - Effective write combining and wear leveling absolutely necessary
- Space will be at a premium
 - Use flash only where necessary
 - Move all home directories / profile data to bulk storage

CUSTOMER EXAMPLE

- Citrix XenDesktop
- 3,000 users
- Vmware ESX 4.0 hypervisor
- HP EVA w/ 250 spindles
- Boot storms degraded performance on entire SAN
- User experience unacceptable

CUSTOMER EXAMPLE (CONT.)

- Migrated write caches to flash array
- Changed write caches to thin provisioned
 - To conserve space
- Boot storms no longer bring down the SAN
- User experience now “better than native”