

# A Brush with Disaster Leads to a Virtualization Based DR Plan at the Las Vegas Valley Water District

Dave Trupkin

Greg Hearn

Sr. Systems Administrators



**VMWORLD 2006**

# Introductions



- Dave Trupkin
  - Sr. Systems Administrator – Intel-based
    - VMware, Antivirus, NAS
- Greg Hearn
  - Sr. Systems Administrator – UNIX-based
    - Backup & Recovery, Oracle, SAN
- Las Vegas Valley Water District
  - Not-for-profit Public Utility
    - Serving over 1 Million People
    - Las Vegas, Blue Diamond, Searchlight, Kyle Canyon, & Jean
  - Microsoft Windows, Sun Solaris, Oracle, NetBackup
    - 200 + Servers in three locations

# Overview



- In the Beginning
- VMware ESX 1.5
- The Disaster
- Planning a DR Strategy
- Implementation
- Daily Operations
  - DR Operations
- Conclusion
  - Lessons Learned
- Q & A

# In the Beginning



## ■ Environment

### > Server and Applications

- AS/400 & OS/2 on Token-ring
- Oracle on Sun Solaris (several stovepipe DB)
- Windows – Each app on its own server
- Development & production apps on same servers

### > Infrastructure

- Old building with many renovations
- Power and cooling
- Floor & rack space

### > Growth led to server sprawl

- Multiple silos of dissimilar hardware
- No consolidated storage (all direct attached)
- Tape backup required same hardware for Windows
- Unreliable Servers

## In the Beginning




- We Invented VMware Server to solve the problem
  - Actually scripting of VMware Workstation
  - IBM introduced ESX Server 1.5 to LVVWD
  - Attended xSeries Conference in Orlando to see ESX
  - Difficult to get evaluation copy without professional services

# VMware ESX 1.5

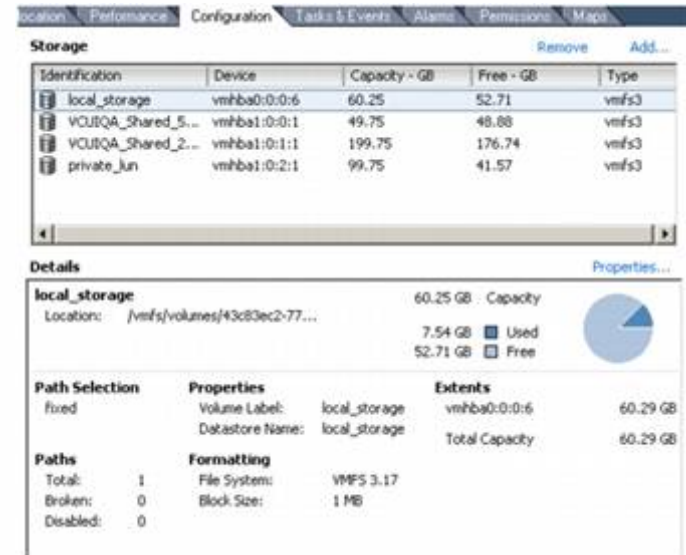


- VMware ESX 1.5
  - ▶ Single IBM x440 4-way
    - No SAN – IBM FastT
  - ▶ Procuring hardware takes time
    - Long lead time for hardware led to developer acceptance
  - ▶ Difficult to gain developer acceptance
    - Every application problem was fault of virtualization
  - ▶ Impossible to gain vendor acceptance
    - Only when there was no other choice



# VMware ESX 1.5

- Upgraded x440 to 8-way/32gb RAM
- Added additional x445 8-way server
  - Acceptance continued slowly
- Introduction of VirtualCenter 1.0
  - Added 9960 & Brocade switches (SAN)
  - Moved virtual servers to SAN
  - Introduced VMotion
- Additions helped with acceptance
  - Additional flexibility
  - Additional functionality
- Still we did not have full acceptance



The screenshot shows the VMware ESX storage configuration interface. It includes a table of storage devices and a detailed view of the 'local\_storage' device.

Identification	Device	Capacity - GB	Free - GB	Type
local_storage	vmhba0:0:0:6	60.25	52.71	vmfs3
VCUBQA_Shared_5...	vmhba1:0:0:1	49.75	48.88	vmfs3
VCUBQA_Shared_2...	vmhba1:0:1:1	199.75	176.74	vmfs3
private_jun	vmhba1:0:2:1	99.75	41.57	vmfs3

**Details**

**local\_storage** 60.25 GB Capacity  
 Location: /vmfs/volumes/43c83ec2-77...  
 7.54 GB Used  
 52.71 GB Free

Path Selection	Properties	Extents
Fixed	Volume Label: local_storage Datastore Name: local_storage	vmhba0:0:0:6 60.29 GB Total Capacity 60.29 GB

Paths	Formatting
Total: 1	File System: VMFS 3.17
Broken: 0	Block Size: 1 MB
Disabled: 0	

# The Disaster



- Dreaded phone call
- Upon arrival some servers still running
  - Oracle and many other redundant services
  - Power temporarily restored (fluctuations)
    - Decision to bring down all remaining servers
- Power surge overloads UPS - Complete power outage
  - Most physical servers down
  - Oracle DB down
  - Did not get time to shutdown VMware or Hitachi
    - Reliability!!
    - All VMs came up clean and in same state
    - AWSOME!





## The Next Day

The District had a power failure in the North Wellfield resulting in the East Administration Building, the Garage, Distribution, & the Modulars not having power.

We are currently working on restoring the AS400 & other Computer Systems for customer Service

We anticipate restoration of power between 9:00 AM & 12 Noon.

for more info - contact Bob Hulshouser  
AT 258-7277

## The Disaster (cont.)

- Limited power restored the following day (24 hours later)
  - Management decision on critical applications
    - Network
    - AS/400 (Customer Billing)
    - Email (Lotus Notes)
    - Oracle (one cluster member)
    - Limited LOB applications
  - Focus was on applications not dependencies
    - Required domain controllers, DNS, etc.
    - More power which we did not have
    - Foresight to have several VM servers



## The Disaster (cont.)

- Brought up selected VMs
  - Support infrastructure
  - Production applications
  - Developer VMs re-tasked
- Ran in this mode for three business days
- High cost of keeping developers idle
  - Additional temporary power feeds installed
  - Full environment restored on temporary power within five business days
- Temporary feeds replaced with permanent power after eight weeks



# Planning a DR Strategy



- Existing Plan
  - Paper only
  - Had no idea on application dependencies
  - Recover physical server infrastructure
- Based on the Tape & Pray concept
  - Find a ballroom
  - Obtain hardware
    - Servers
    - Tape libraries
    - Etc...
- 7-10 days estimated
  - In an ideal world!

# Planning a DR Strategy



- New Plan
- Application administrators and developers
  - Identified application dependencies
- Appreciation for VM efficiencies and performance
  - VMware provides missing step
  - ESX infrastructure expanded
    - Now includes 15 hosts (standalone and BladeCenter)
    - New servers must be virtual unless impossible
- District priorities shift
  - Prior emphasis solely on water distribution
  - New appreciation for ancillary services
  - Expectations on LOB applications
- May not have existing facility next time
  - Fortunate that we did not lose the facility

## Implementation

- Leased space at co-location facility
- Purchased second Hitachi
  - Upgraded 9960 with Sun 9990
  - Proved TrueCopy replication
- Purchased dark fiber
  - Two dedicated pairs (9.6 Miles)
    - One for network & One for SAN
  - Existing SONET OC-48
    - Added OC-12 prevision
- Installed CWDM on pair for SAN
  - Increased our data path to 4 x 2GB
- Additional Sun 9990 hardware & software
  - Cache & TrueCopy





# Implementation

- ESX Server hosts
  - 1 xSeries 445 (8-way 32GB)
  - 1 xSeries 460 (4-way 32GB)
  - 1 BladeCenter chasis
- Backup and Recovery services
  - Virtual Tape Library (VTL)
  - NetBackup Media/master
- SunFire 6800
  - 3 domains (8-way 32GB)
  - Oracle database
  - Capacity for production applications



# Data / LUN Configuration



- In the beginning
  - Large LUNS
  - Concatenation
- LUN contention
  - File Locks
- Created smaller LUNs
  - Different sizes
  - No concatenation
- Create logical LUN groups
  - Manageability
  - Data integrity
  - Based on data/server layout



# Data Replication

- Use of Synchronous & Asynchronous
  - Oracle Archives Synchronous
    - Never split the pairs
    - High change rate
  - Oracle Database (RAC)
    - Asynchronous with splits
    - Medium change rate
  - VMware servers
    - Asynchronous with splits
    - About 5 TB total in 8 groups
    - Medium change rate
- All replication continues 24 x 365



# Data Replication

## ■ Hardware

- Sun Hitachi 9990 (USP) to HDS 9960
  - Higher cache install on 9990
- ADVA FSP 2000 CWDM/DWDM
  - 8 SFP (4 network & 4 SAN)
- Brocade SilkWorm 4100 Switches



## ■ Software

- Hitachi TrueCopy
  - One primary to one secondary (1 PVOL to 1 SVOL)
  - Need two ports for RCU and MCU on both arrays
- Hitachi HORCM CLI
  - Not difficult to use, but difficult to troubleshoot
  - Scripts & cron jobs

# TrueCopy Configuration



- Synchronous copy of Oracle archive logs
  - 100 – 200 GB
  - High change rate
  - Impacts performance on production
    - Requires additional cache to keep performance up
- Asynchronous copy of Oracle DB
  - 1 TB RAW or 720 GB of data
  - Split pairs at 06:00, 12:00, 18:00 & 24:00 hrs
  - Minimal impact in performance on production LUNs
- Asynchronous copy of VMware VM servers
  - 5 TB RAW or 4 TB data
  - Refresh in groups starting at 18:00
  - No impact on performance on production

# Data Replication Breakdown




Asynchronous = 6.5 TB

- VMware VM servers
- Oracle DB
- Other data

Synchronous = 200 GB


- Oracle Archives



# Data Replication Breakdown



### Duration



# Conclusion

- Lessons learned
  - Disaster comes in many forms and sizes
  - Regular review needed of application priorities
  - Application dependencies
  - May need to setup rest areas for employees
  - Critical knowledge may be in someone else's head
  - Communication is critical
    - Co-workers
    - End users
  - Synchronous is not always better than asynchronous



# Conclusion



- DR site adds value to daily operations
  - Protection from failed patches
  - Protection from failed upgrades
  - Provides “seed” data for populating development environments
  - Peace of mind (good for sleeping)

# Questions ???



## Presentation Download

Please remember to complete your  
**session evaluation form**  
and return it to the room monitors  
as you exit the session

The presentation for this session can be downloaded at  
**<http://www.vmware.com/vmtn/vmworld/sessions/>**

Enter the following to download (case-sensitive):

**Username: cbv\_rep**  
**Password: cbvfor9v9r**

**VMWORLD 2006**

