# Migrating Server Operations from Remote Sites to the Datacenter for Disaster Recovery and Protection

Angelo Armenteros – Director of Network Systems

Eddy Caballero – Storage Analyst

Greenberg Traurig

# Agenda

- Introduction
- Company Profile
- DR Challenges
- Previous DR environment
- Design Goals
- Implementation
- Implementation Issues
- Conclusions
- Q & A

# Company Profile

- Greenberg Traurig


- International Law Firm – 29 offices in US, Europe, Asia
- 100+ Practice areas: Appellate, Aviation, Corporate, Entertainment, Environmental, Governmental, Healthcare, Intellectual Property, Labor, Litigation, Real Estate, Securities,  Tax, T&E, etc….
- 1500+ attorneys, 3500+ employees
- Ranked No. 1 in the U.S. for 5-year growth leaders in The National Law Journal Millennium NLJ 250 annual survey of the nation's 250 largest law firms
- Ranked No. 7 on The American Lawyer's 2006 Am Law 100 listing of the largest law firms in the U.S., based on number of lawyers

# DR Challenges

- Design constraints
  - Highly distributed environment
    - Many remote offices
    - Each office designed to operate independently
  - Two data centers
  - Wide variety of WAN links
    - T1's to 100Mb
  - Mixture of transactional and non-transactional data
    - MS SQL
    - MS Exchange
    - Windows based file servers
    - Domain controllers
    - Other miscellaneous servers

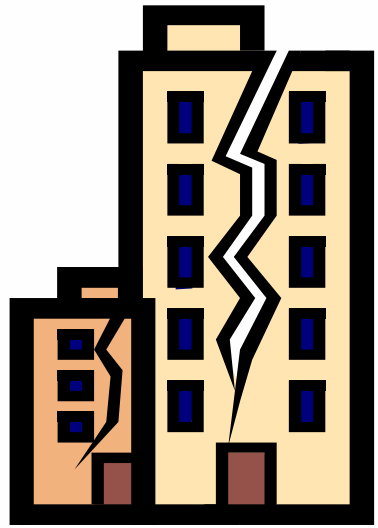

# Previous DR Environment

- Software (host-based) replication + VMware ESX
- Issues:
  - Required loading software into each guest OS
    - Consumed guest resources
    - Potential conflicts with other software
      - Backup agents, anti-virus, and other modules
  - Could not easily replicate entire guest as a bootable image
  - Excessive re-mirror events
    - Certain conditions required replicated data to be re-synced
  - Fail-over / Fail-back required extensive tinkering with replicated guest
    - Machine names, DNS registrations, IP addresses and more
  - Fail-over requires “standby” host (physical or virtual)
  - Limitations of source-target configurations
    - One-to-one vs. many-to-one

# Design Goals

- Quick fail-over / fail-back
- Transparent fail-over / fail-back
  - No changes required to replicated servers
  - No changes required to client devices
- Replicate each server as a bootable unit
  - Fast “cold migrate” functionality
- Storage-level replication
  - No OS involvement
- Leverage capabilities provided by VMware
  - Data encapsulation
  - Hardware abstraction
- Transactionally consistent data (vs. crash-consistent)
  - Especially important for MS Exchange and MS SQL

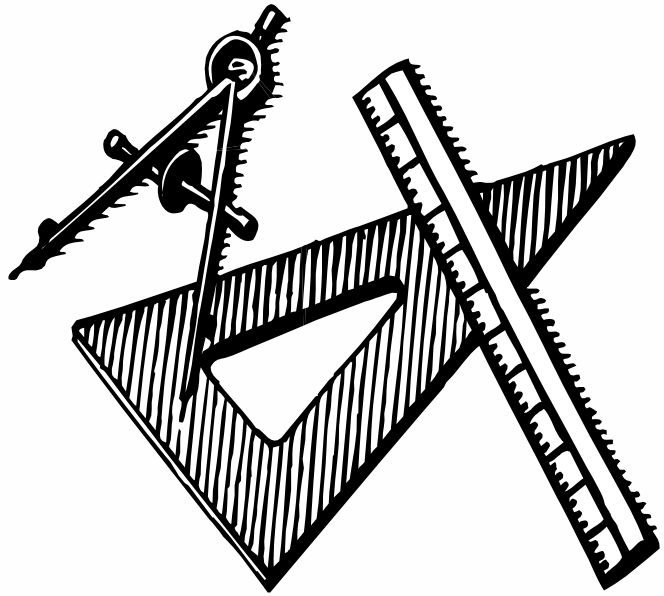

# Implementation

- Platform Components:
- VMware ESX 2.5.x / VC 1.x
- Storage, virtualization, management and replication
  - IPStor 5.x from FalconStor
  - SCSI or FC (SAN) shared storage array
- Networking
  - Layer 3 routing switches at each network core
- Scripting and integration
  - Tcl/Tk and Expect for scripting and integration
- Remote access
  - Citrix Presentation Server 4
  - MS Outlook Web Access



CITRIX

# Implementation

- VMware ESX 2.x / VC 1.x
- Each remote office has full complement of servers required to work independently
  - DC, SQL, Exchange, file server, print server, document management, web proxy, SMS  and other miscellaneous servers
  - Typical deployment is around 10 – 12 virtual servers per site
  - 2 – 5 physical ESX hosts per site
- Not all servers required for DR
  - Only SQL, Exchange, file server and DMS related are critical

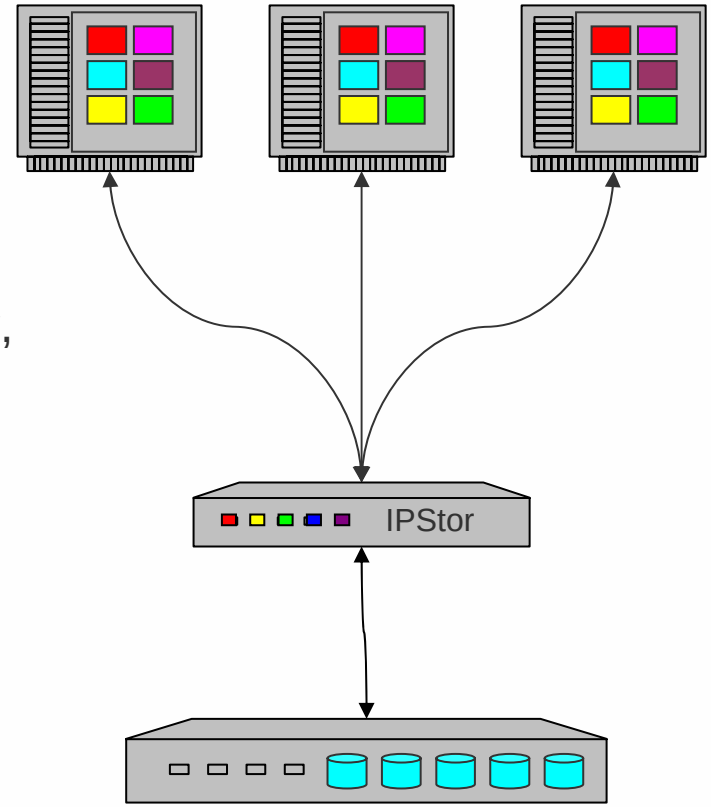

# Implementation

- Storage
- Fiber Channel SAN or SCSI-based shared storage
- FalconStor IPStor
  - Storage presentation
  - Storage virtualization
  - Snapshots / Mirroring
  - Replication
- Replication
- CDP vs. Snapshots
- RPO vs. Transactional consistency

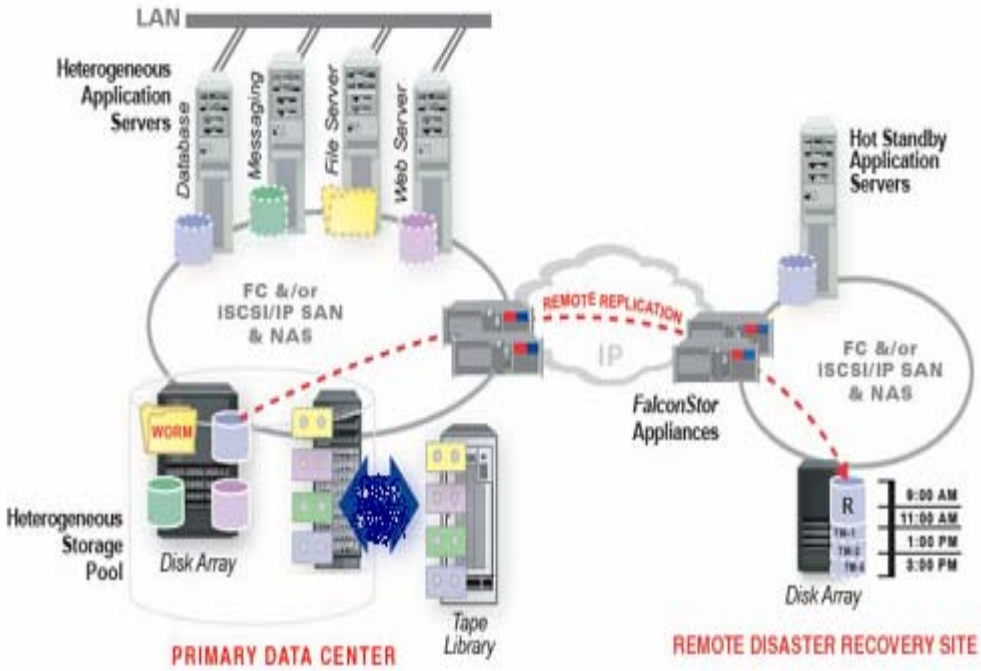

# Implementation

- Storage Virtualization
- Provide consistent storage presentation and management regardless of underlying storage type
  - SCSI, SATA, SAN, etc.
- Easily migrate between physical storage systems
- Add functionality to existing storage
  - Snapshots / Cloning
  - Replication / Mirroring
    - Synchronous
    - Asynchronous
    - Continuous
    - Periodic

# Implementation

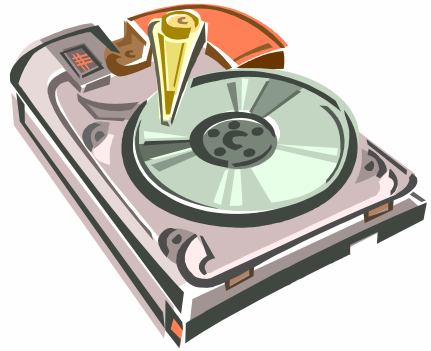

- RDM vs. VMFS
- RDMs are more difficult to manage, but….
  - Much more practical in SAN environment where snapshots are used
- VMFS
  - Snaps on VMFS track I/O for all VMDKs, not just the one desired for the snap
  - Snaps can only be presented back to ESX hosts
- RDM
  - Snaps of RDMs only track I/O for one specific LUN/volume/drive
  - Snap can be presented to ESX or to a physical host
- ESX needs better RDM management
  - Need method to globally ID, track and manage an RDM independent of ESX host it was created on

# Implementation

- Networking
- Layer 4 Routing switch at network core
- Each office subnetted into multiple subnets / VLANs
- All servers on single, dedicated subnet / VLAN
- Dynamic routing protocol (RIP v2)
- VLAN / subnet transportable throughout network

- Example:
  - 10.1.1.x – Servers
  - 10.1.2.x – Printers
  - 10.1.3.x – User 1
  - 10.1.4.x – User 2
  - …..

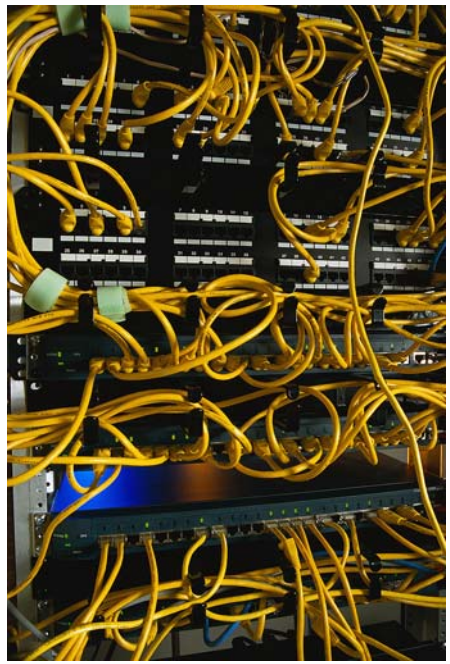

# Implementation

- Scripting and Integration
  - Used to facilitate talking to a number of dissimilar systems
    - ESX server
    - IPStor
    - Routing switches
  - Tck/Tk with Expect
    - Easiest method to automate CLI interfaces

TCL
radically simple

# Implementation

- System Access after fail-over
- LAN/WAN access
  - No changes needed to existing systems
  - All failed-over servers have same names and IP addresses
  - Only change was network route
- Remote Access
  - Citrix Secure Gateway
  - Citrix Presentation Server 4
  - MS Outlook web access
  - Limited VPN access

# Implementation

Failover process

- Shutdown source VMs
- Shutdown source ESX servers
- Flush any pending replication data
- Shutdown source router VLAN interface

(Planned)

---

(Planned or Unplanned)

- Activate target router VLAN interface
- Present replicated data to fail-over ESX hosts
- Perform any VM guest setting adjustments needed
  - RDM presentation, VMX tweaks, etc…
- Boot DR VM guest OS

# Implementation Issues

- Replication
  - Know your data change rates
  - Identify and separate critical vs. non-critical data
  - WAN capacity
  - WAN latency
  - Data compressibility
- Storage capacity and I/O bandwidth
  - Space for snapshots and replicated data
  - Snapshots and/or CDP require extra I/O bandwidth
  - In fail-over mode, extra storage capacity needed (potentially) for efficient fail-back
  - Use RDM's to isolate snapshot I/O

# Implementation Issues

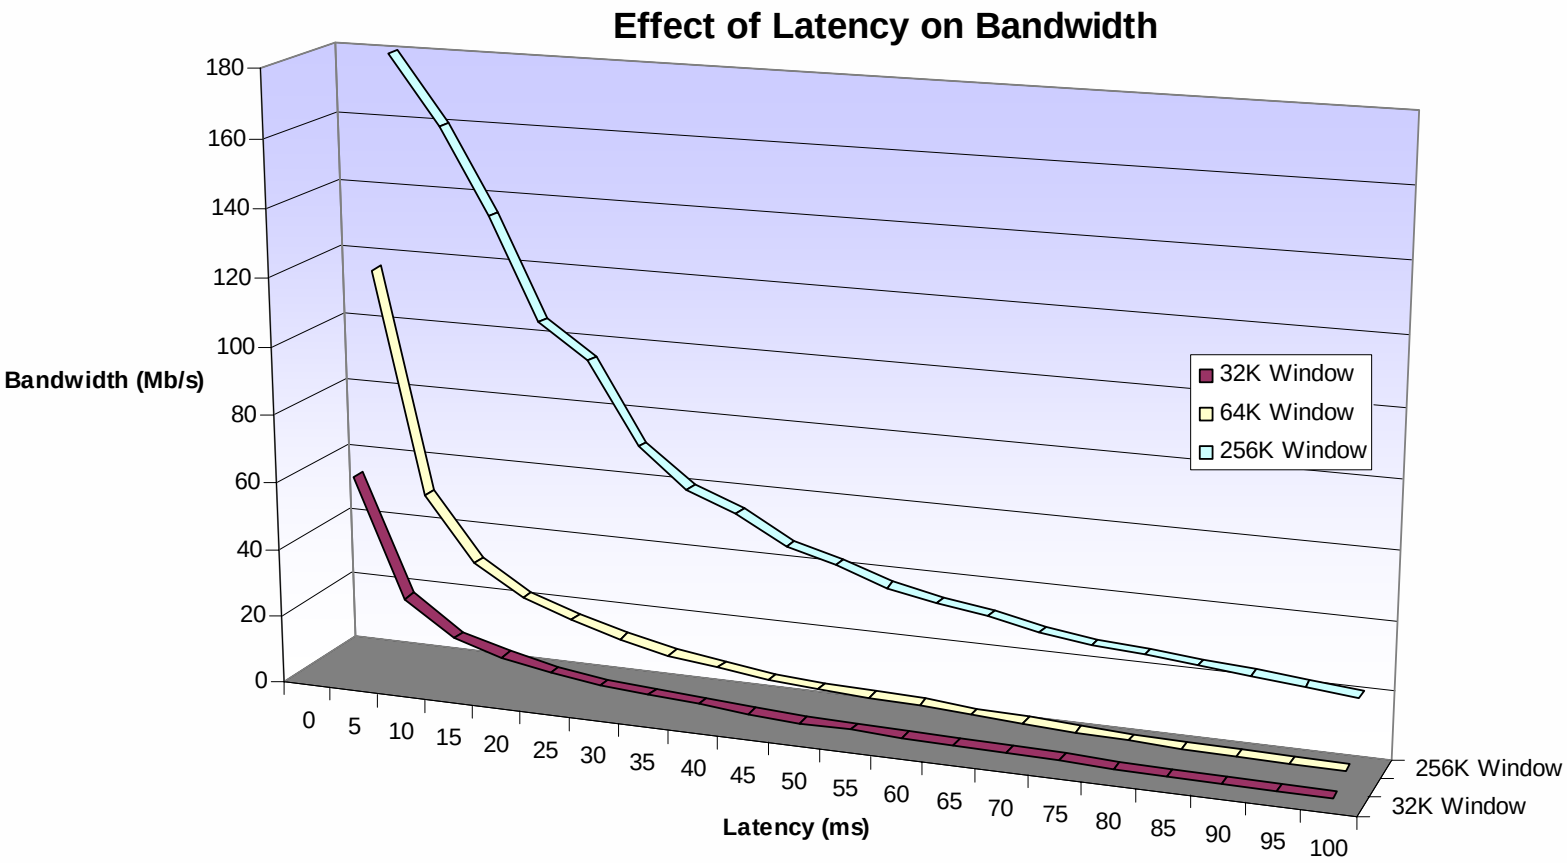

# Conclusions

- Storage virtualization together with VMware greatly facilitates DR replication and fail-over
- Storage level replication solves many replication issues
  - No resource utilization on replicated server
  - Replication not affected by server OS issues
  - Except (nothing is perfect!)
    - Disk defragmentation
    - Easy file/folder exclusion
- Entire subnet fail-over eliminates need to change server and client device settings
  - Single server can be failed-over with a little more effort

# Conclusions

- Uses:
  - DR
  - Remote office maintenance
  - Upgrades
  - Office moves
  - Testing

# Q & A

# 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