November 2011

Disaster Recovery with VMware Site Recovery Manager and Exchange 2010 High Availability in a VMware Virtual Environment

Yury Magalif, MASE, VCP

Senior Systems Architect

ymagalif@promedianj.com

PROMEDIA

TECHNOLOGY SERVICES, INC.

Thank you!

- I would like to thank my colleagues Peter Alberto and Ralph Carter for valuable tips in the creation of this presentation.

Agenda

I. Basics of HA and DR.

II. Why Exchange native HA/DR?

III. Sizing for Exchange.

IV. Building tips for Exchange.

V. Why SRM native DR?

VI. Building tips for SRM.

VII. Case Study

I. Basics of High Availability & Disaster Recovery

Challenges of Traditional Disaster Recovery

Expensive


Software


Hosts


Storage


Facilities

>$10K per app

Complex Recovery Plans

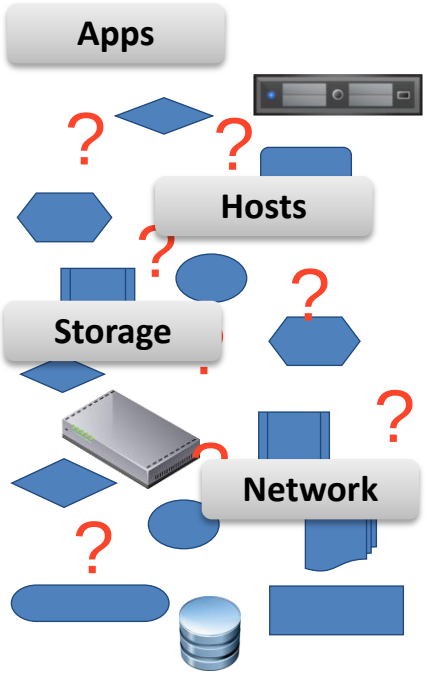

Unreliable Failovers

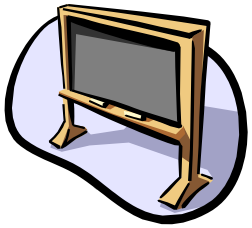

Failure to meet business requirements

- Recovery takes days to weeks
- Too much time and resources consumed

Why Virtual DR?

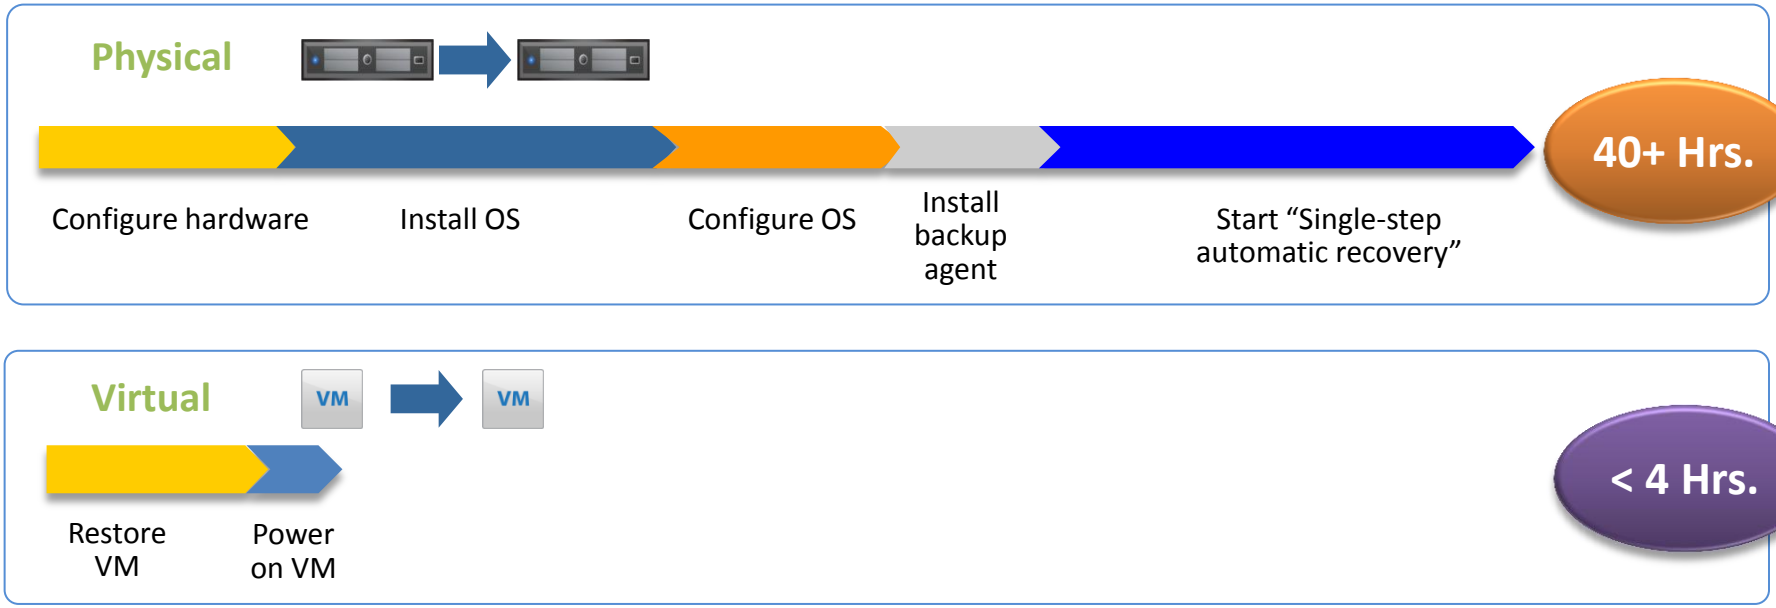

Simplify recovery

- No operating system re-install or bare-metal recovery
- No time spent reconfiguring hardware

Standardize recovery process

- Consistent process independent of applications, operating systems and hardware

Differences between High Availability & Disaster Recovery

- How big is your Problem?
- Flat Tire, Windows won't boot, One server died.
- High Availability (HA) = Tire Sealant, Run-Flat Tire
- Shredded Tire, Hurricane Irene flood, No power
- Disaster Recovery (DR) = Spare Tire


The Michelin PAX System

Basics of RTO and RPO

- Recovery Point Objective (RPO) = if your tire was shredded, can you lose your bumper?
- IT: User emails for the past 2 hours were destroyed, but it's OK
- Recovery Time Objective (RTO) = How fast can you change shredded to spare tire?
- IT: it will take me 1 hour to recover the email server so user can send email.

II. Why Exchange 2010 HA and Load Balancing?

- VMware HA – already included if you have a VMware cluster, but 2-5 min RTO (boot time of virtual machines)
- Double-Take from Vision Solutions, Neverfail – great solutions, RTO in seconds, but third party costs
- Exchange 2010 HA for Mailbox Role and Network Load Balancing (NLB) for CAS/Hub role
  - Native HA, DR, and Load balancing from Microsoft
  - RTO – 5-30 seconds for HA, ~30 minutes for DR
  - Comes free with Exchange
- Hardware Load Balancers like F5, Cisco ACE for CAS/Hub -- Great boxes but have a cost

III. Exchange 2010 Sizing 01

1. Use Exchange 2010 Mailbox Server Role Requirements Calculator
2. Discovery – run AD Topology Diagrammer, Exchange Profile Analyzer, Exchange Pre-Deployment Analyzer, Exchange Best Practices Analyzer.
3. Possibly run Sydi for Exchange, sydiproject.com

Exchange 2010 Sizing 02

- Size Virtual Exchange servers per Microsoft formulas for physical servers.
- To get local HA and remote DR for the mailbox role, build at least 3 Exchange DAG servers.
- To get local HA and remote DR for the CAS/Hub roles, build at least 3 NLB servers.
- For NLB to work, use a stretched VLAN. NLB is not possible across subnets.

Exchange 2010 server role	Minimum supported	Recommended maximum
Edge Transport	4 GB	1 GB per core (4 GB minimum)
Hub Transport	4 GB	1 GB per core (4 GB minimum)
Client Access	4 GB	2 GB per core (8 GB minimum)
Unified Messaging	4 GB	2 GB per core (4 GB minimum)
Mailbox	4 GB	4 GB plus 3-30 MB additional memory per mailbox: The total required memory is based on the user profile and database cache size. For more information about how to determine the total required memory, see Understanding the Mailbox Database Cache.
Client Access/Hub Transport combined role (Client Access and Hub Transport server roles running on the same physical server)	4 GB	2 GB per core (8 GB minimum)
Multiple roles (combinations of Hub Transport, Client Access, and Mailbox server	8 GB	4 GB plus 3-30 MB additional memory per mailbox: The total required memory is based on the user profile and database cache size. For more information about how to

Mailbox Role HA – Database Availability Group (DAG)

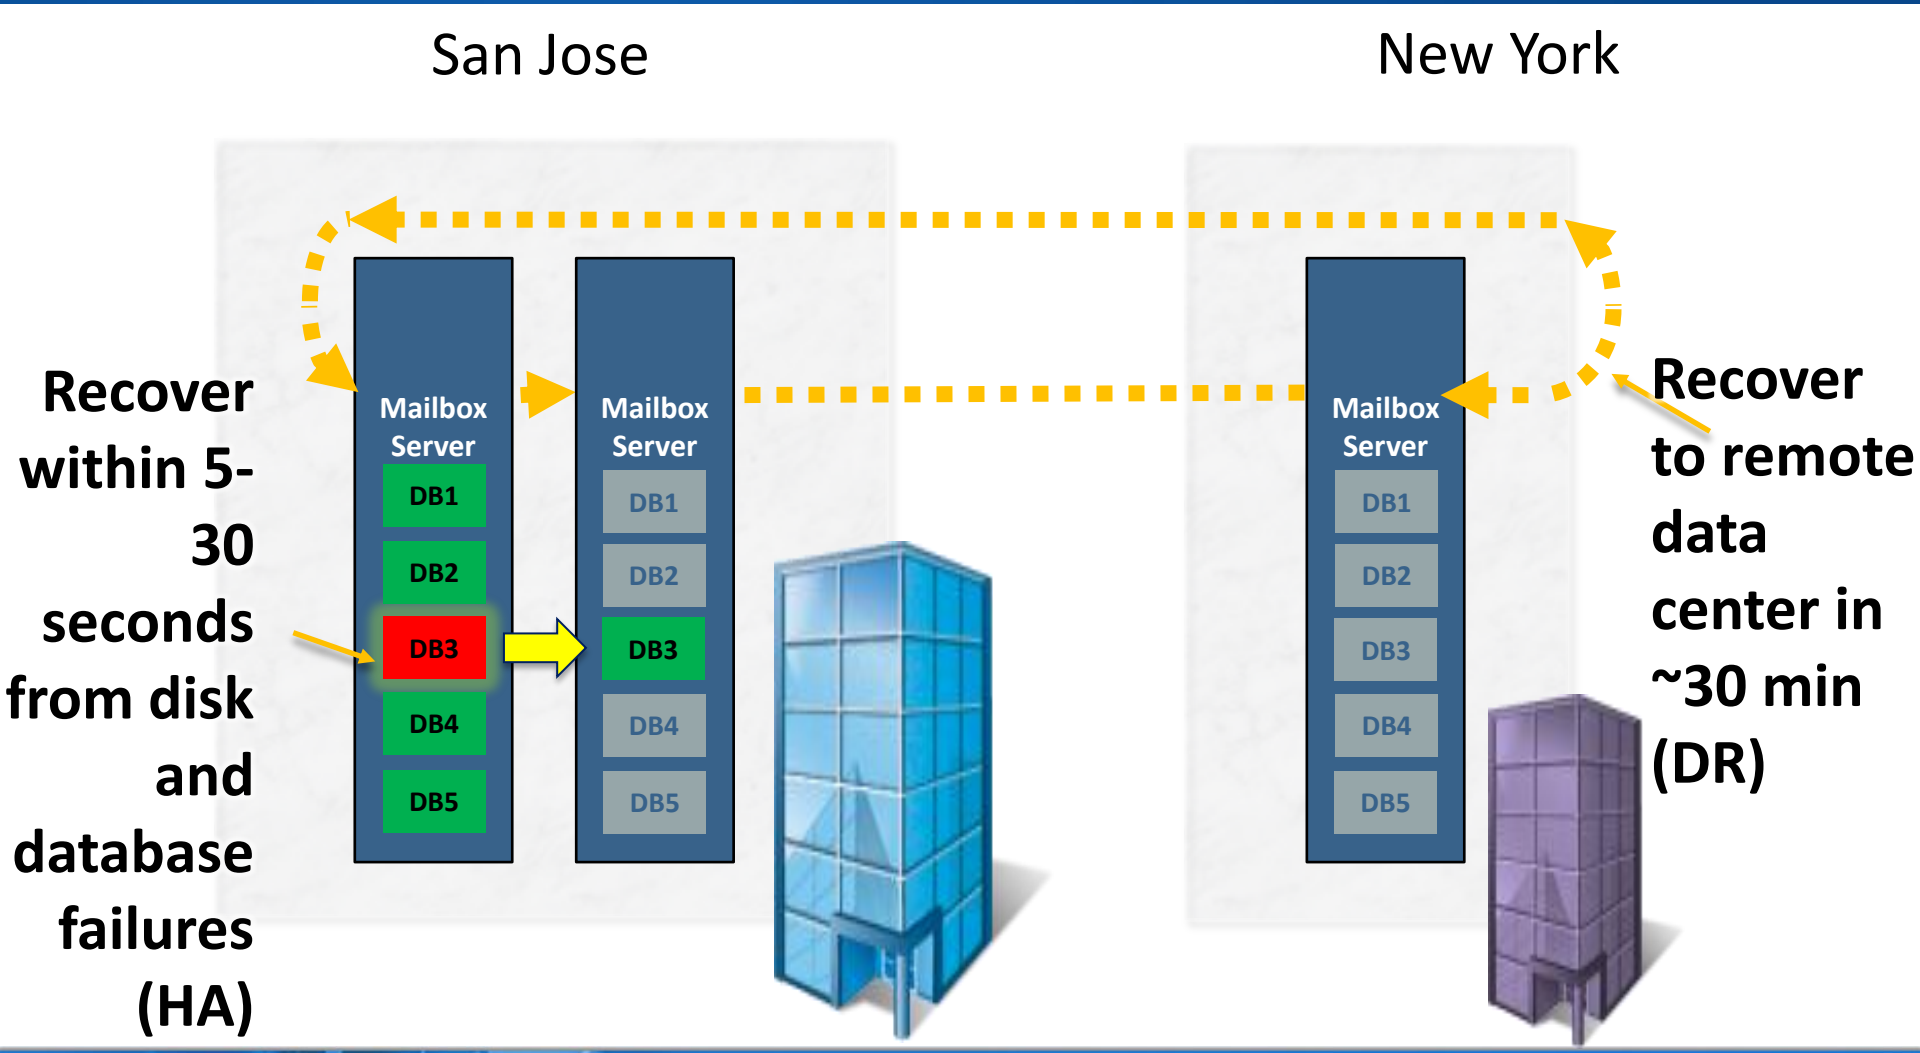

CAS/Hub Role HA – Network Load Balancing (NLB)

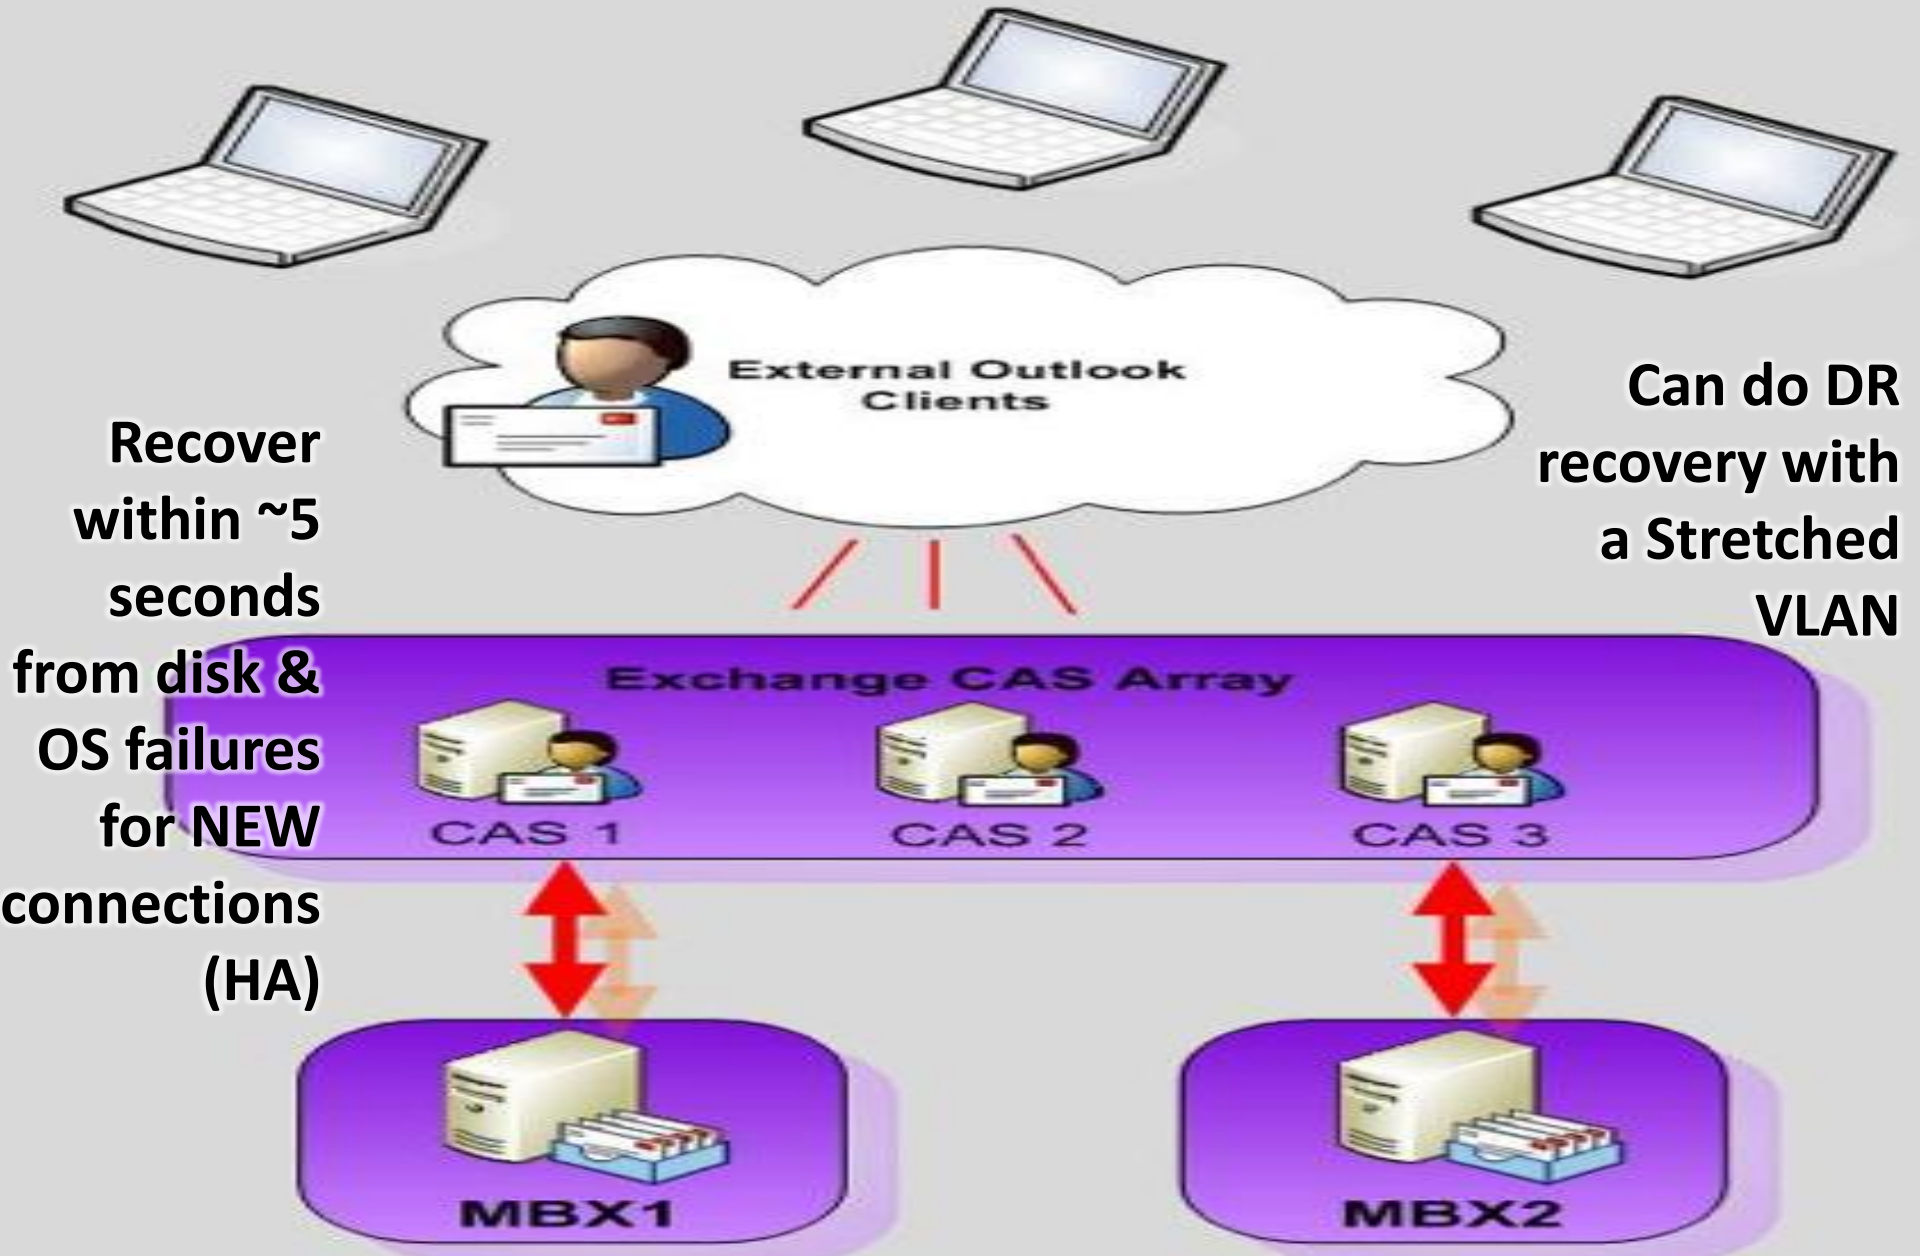

Recover within ~5 seconds from disk & OS failures for NEW connections (HA)

Can do DR recovery with a Stretched VLAN

IV. Exchange 2010 Building Tips

- Enable Datacenter Activation Coordination Mode
- Use Multicast for NLB servers
- Make sure to have an odd number of DAG nodes, or even number and file share witness
- If you have an archiving system like Symantec Enterprtise Vault, point it at the CAS servers.
- Do not turn off IPv6 in a clustered environment or when using DAGs because Windows Server 2008 R2 Clustering uses IPv6 for internal communication.
- Never separate mailbox and Client Access Servers with a firewall -- keep them in the same network, or use MS Forefront TMG firewall in the DMZ

Why Not Site Recovery Manager for Exchange?

RTO: 30 minutes to hours

RPO: Flexible based on storage replication

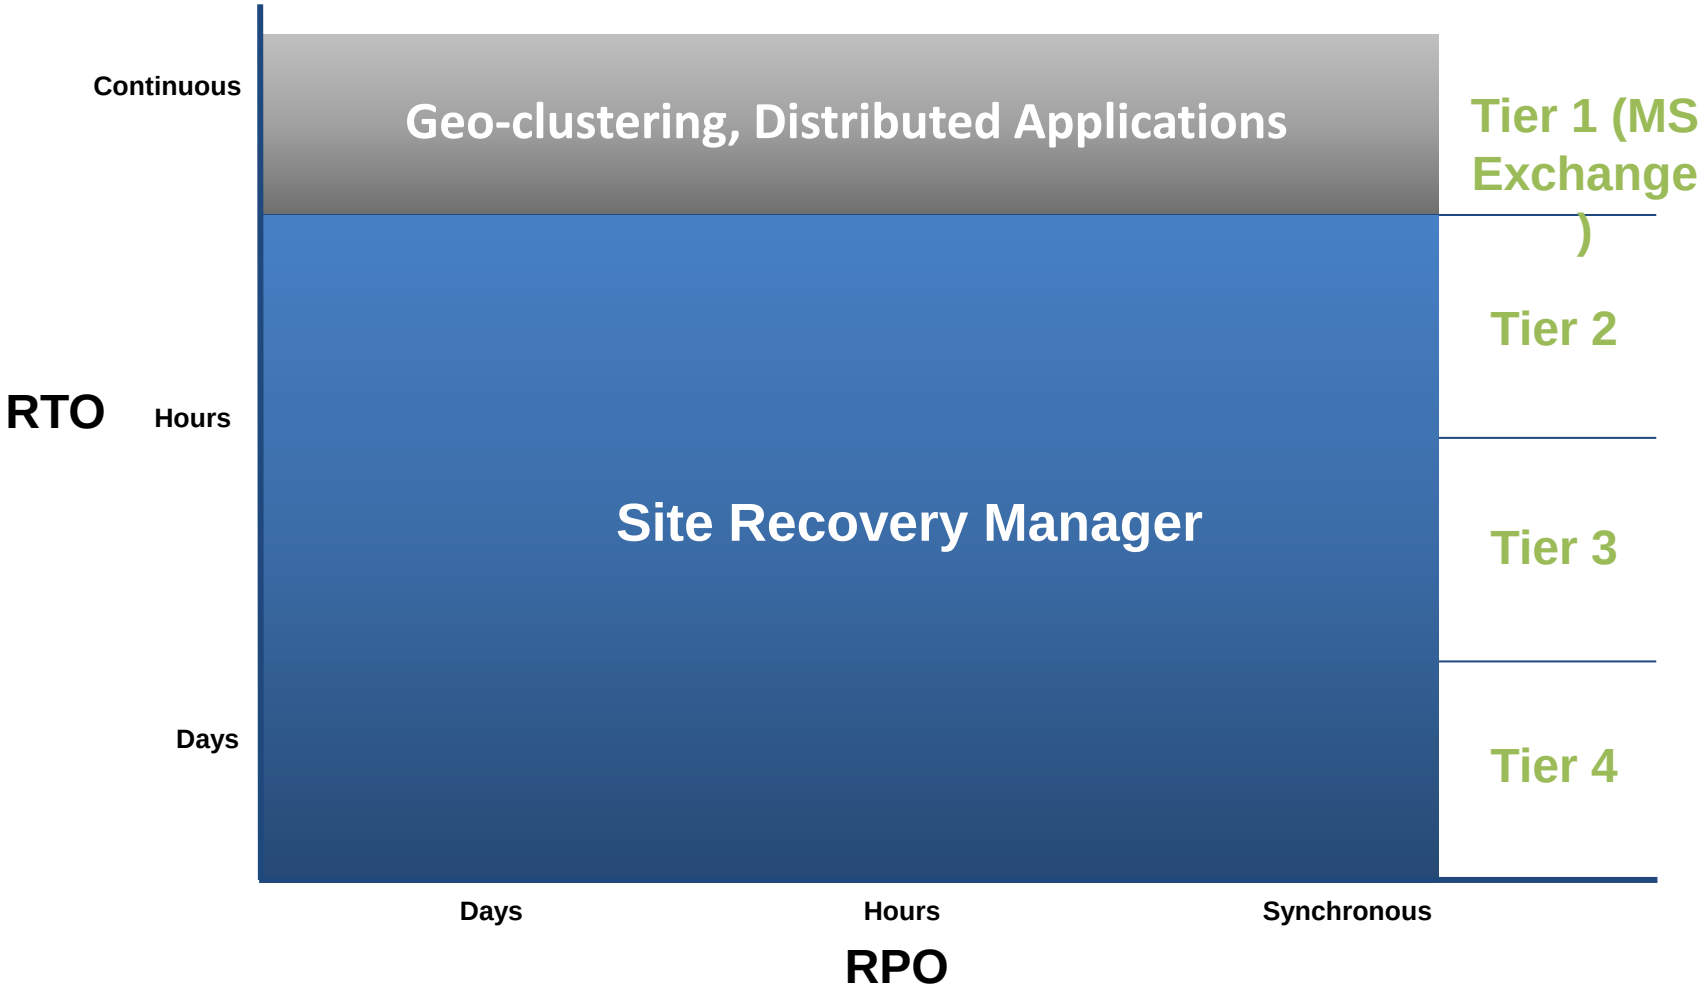

V. Why VMware Site Recovery Manager?

- Quest vReplicator, Double-Take – good, but not native and cannot synchronize with storage array replication
- Site Recovery Manager 5
  - Native DR from VMware
  - The only solution that automates DR with replication between hardware storage arrays.
  - RTO – 30-60 minutes for DR, depending on # of VMs.
  - Can do software replication if you don't have storage array replication.
  - Can test DR without impact to production.
  - Has automated fail back (not a 1 button operation, takes 3 steps)

VMware Site Recovery Manager (DR)

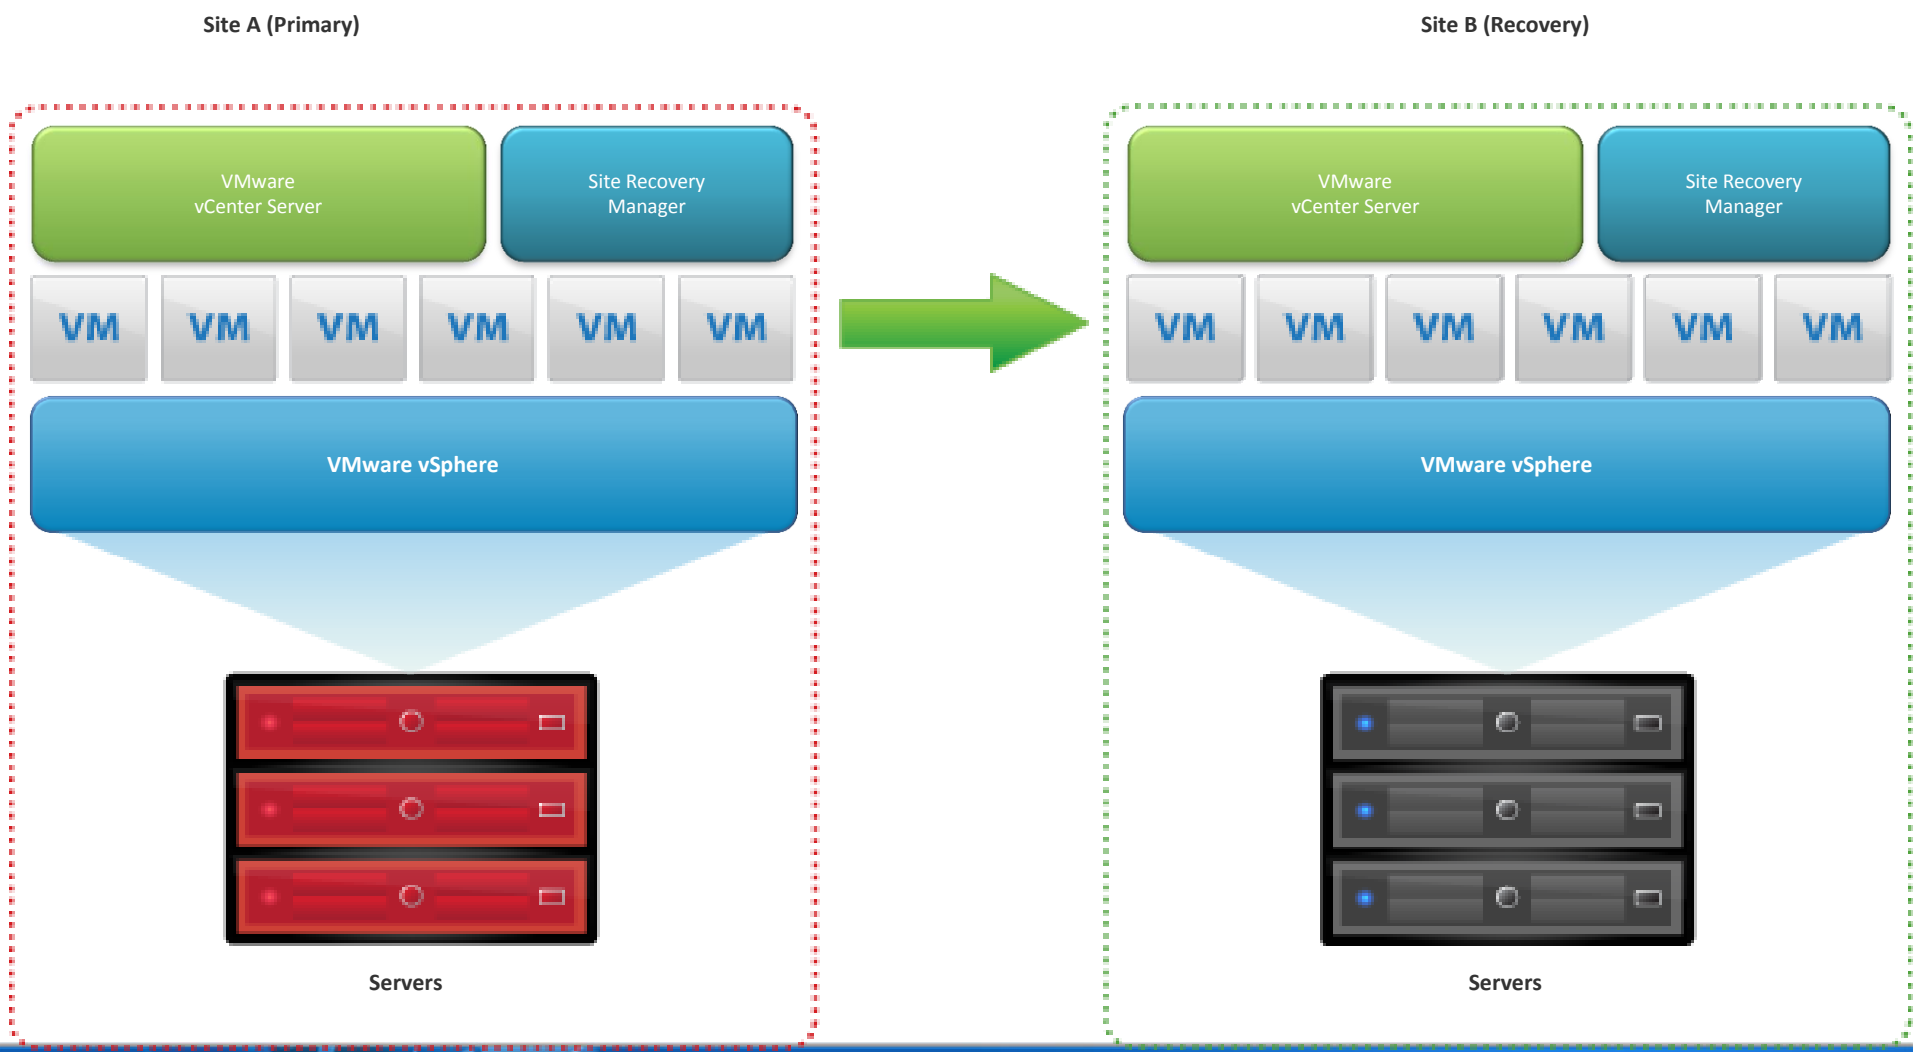

SRM Provides Choice of Replication Options

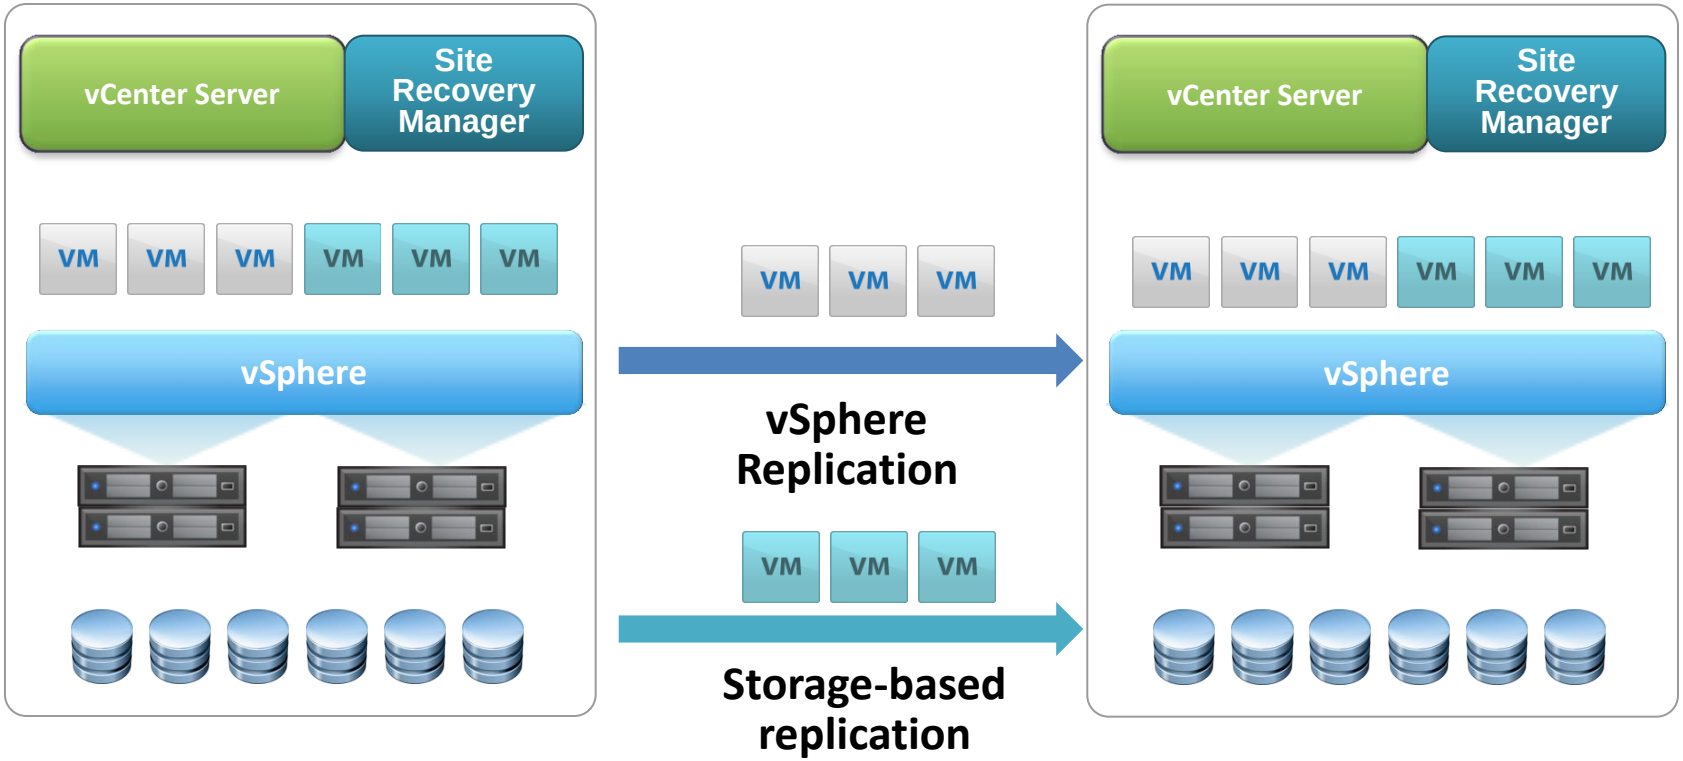

vSphere Replication
Simple, cost-efficient replication for Tier 2 applications and smaller sites

Storage-based Replication
High-performance replication for business-critical applications in larger sites

vSphere Replication Complements

	Replication Provider	Cost	Management	Performance
vSphere Replication	VMware	 • Low-end storage supported • No additional replication software	 • VM' granularity • Managed directly in vCenter	 • 15 min RPOs • Scales to 500 VMs • Application consistency for planned migrations only
Storage-based Replication	EMC², NetApp, HP invent, 3PAR, compellent, FUJITSU, LSI, DELL EQUALLOGIC, Xiotech, FalconStor Software, IBM, Hitachi Data Systems	 • Higher-end replicating storage • Additional replication software	 • LUN – VM layout • Storage team coordination	 • Synchronous replication • High data volumes • Application consistency possible

Automate DR Failover & Migration Processes

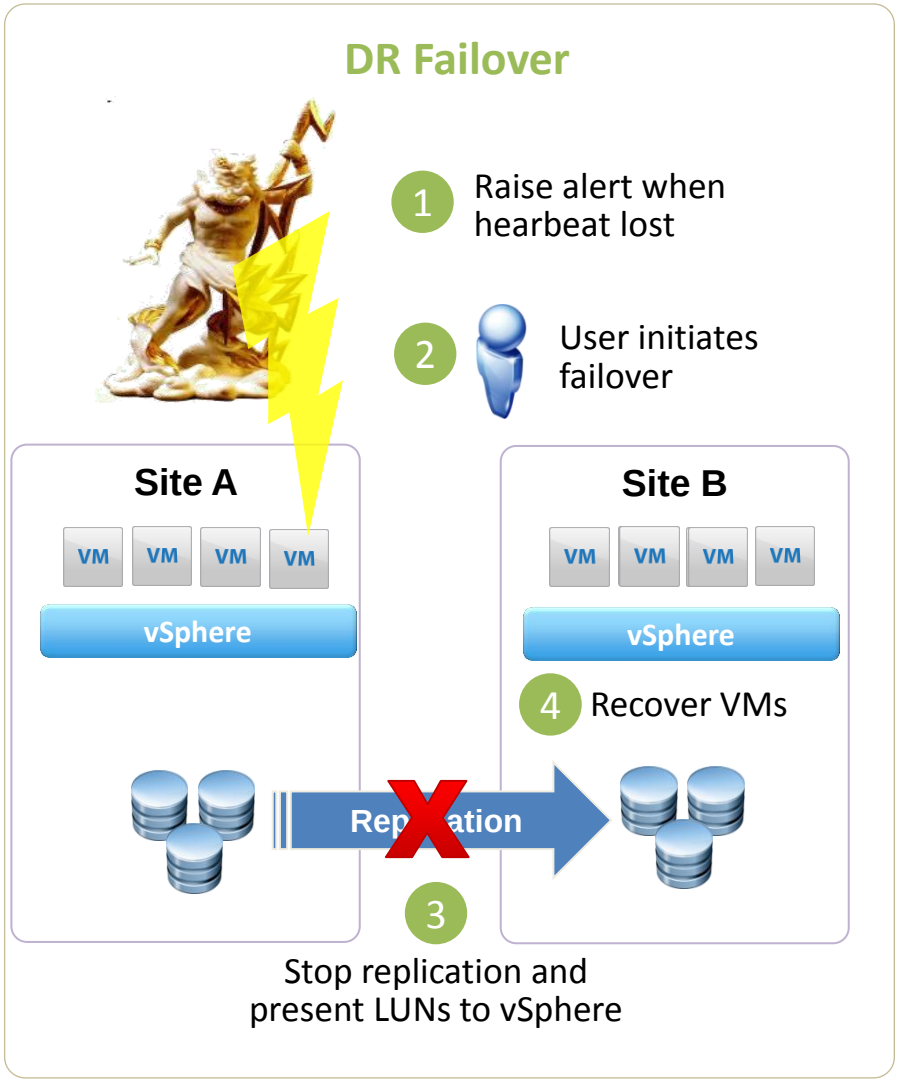

Overview

Automatically detect site failures

- Require user to manually initiate failover

Automate recovery process

- Stop replication and present replicated LUNs to vSphere
- Execute user-defined recovery plan

Benefits

Ensure fast and predictable failovers and migrations

- Consistently meet business requirements

Minimize risk of user errors

Automated Failback

Automated Failback

Reverse original recovery plan

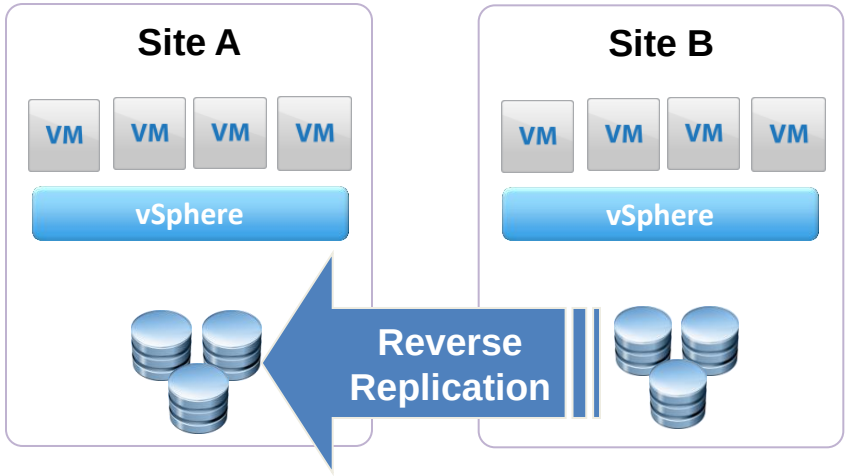

Overview

Re-protect VMs from Site B to Site A

- Reverse replication
- Apply reverse resource mapping

Automate failover from Site B to Site A

- Reverse original recovery plan

Restrictions

- Does not apply if Site A has undergone major changes / been rebuilt
- Not available with vSphere Replication

Benefits

Simplify failback process

- Automate replication management
- Eliminate need to set up new recovery plan

Streamline frequent bi-directional migarations

Beyond DR: Preventive Failovers

Unplanned Failover

Recover from unexpected site failure
- Full or partial site failure

The most critical but least frequent use-case
- Unexpected site failures do not happen often
- When they do, fast recovery is critical to the business

Preventive Failover

Anticipate potential datacenter outages
- For example: in case of planned hurricane, floods, forced evacuation, etc.

Initiate preventive failover for smooth migration
- Leverage SRM 'planned migration' to ensure no data-loss
- Automated Failback enables easy return to original site

Planned Migration

Most frequent SRM use case
- Planned datacenter maintenance
- Global load balancing

Ensure smooth migrations across sites
- Test to minimize risk
- Execute partial failovers
- Leverage SRM 'planned migration' to ensure no data-loss
- Automated Failback enables bi-directional migrations

VI. SRM Building Tips

- Use either FQDN or IP addresses, but do NOT mix
- Do not expect the arrays to pair right away – give it time
- Stretched VLAN is best, otherwise use IP customization
- With a stretched VLAN, use
- When doing Failover tests, you can make a self-enclosed VLAN to test actual users
- For Stretched VLAN use Hot Standby Router Protocol (HSRP) from Cisco to balance gateway address across sites
- Test Failover often, and work on fixing errors.

VII. Case Study – Ridgewood Public Schools 01

- Blue Ribbon School District
- Kindergarten through 12th grade
- 10 buildings
- 5600 students, 400 faculty
- Fiber Ring connecting all schools
- Redundant providers (Verizon & Lightpath)

Case Study – Ridgewood Public Schools 02

- Disaster types:
  - Flooding in buildings
  - Power Outages
  - Network outages
  - Affected by recent snowstorm & Hurricane Irene
- Critical Infrastructure
  - MS Exchange
  - Student Information System
  - Shortel Voice Servers
  - Active Directory
  - Financial System

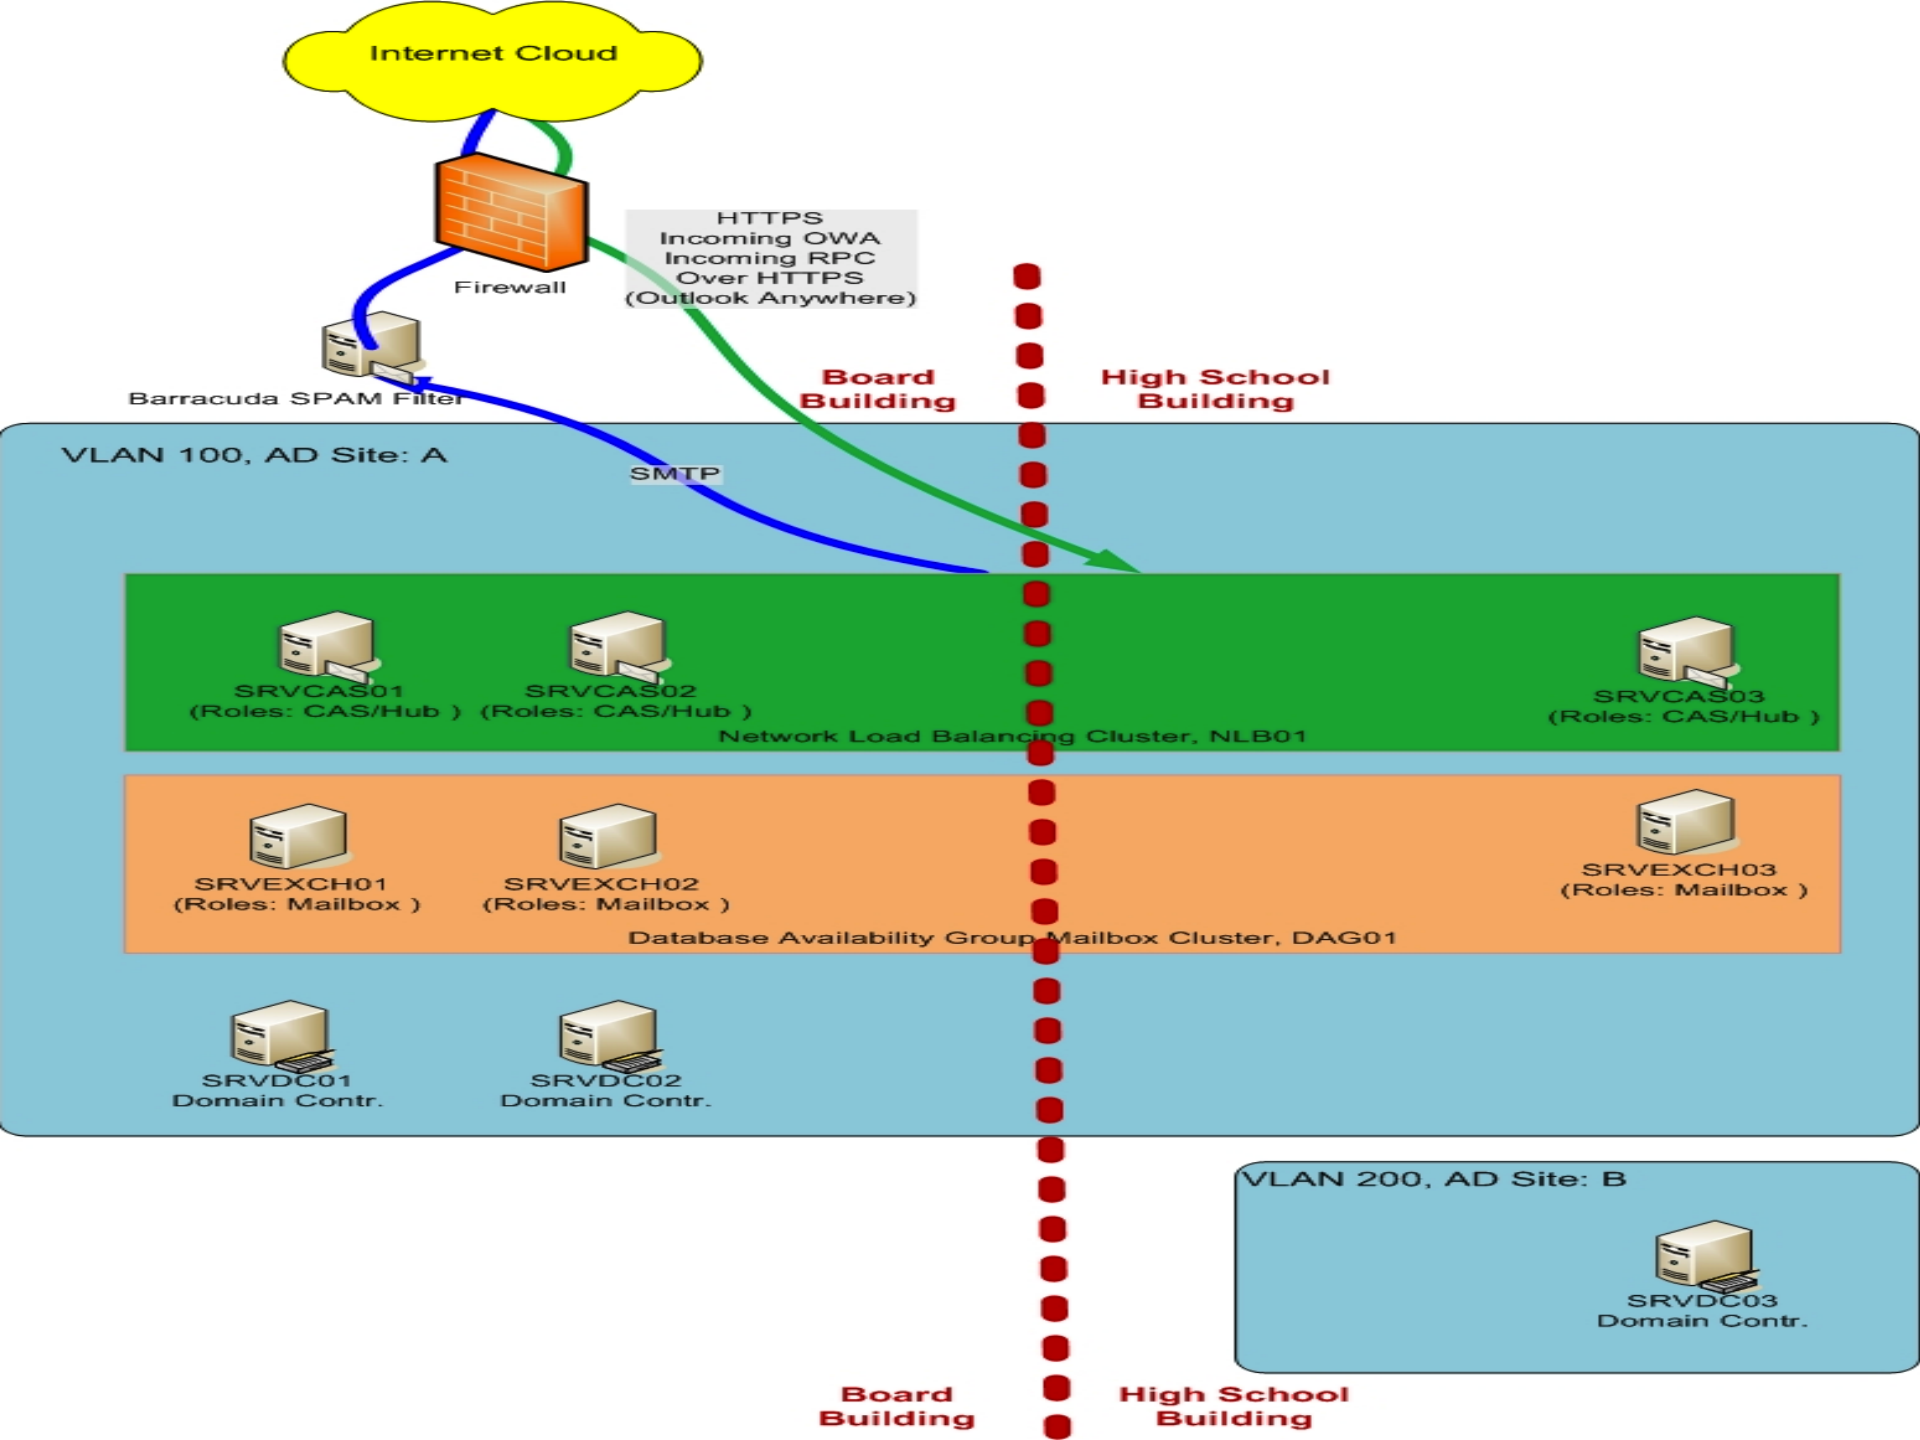
Internet Cloud
Firewall
HTTPS
Incoming OWA
Incoming RPC
Over HTTPS
(Outlook Anywhere)
Barracuda SPAM Filter
Board Building
High School Building
VLAN 100, AD Site: A
SMTP
SRVCAS01
(Roles: CAS/Hub)
SRVCAS02
(Roles: CAS/Hub)
SRVCAS03
(Roles: CAS/Hub)
Network Load Balancing Cluster, NLB01
SRVEXCH01
(Roles: Mailbox)
SRVEXCH02
(Roles: Mailbox)
SRVEXCH03
(Roles: Mailbox)
Database Availability Group Mailbox Cluster, DAG01
SRVDC01
Domain Contr.
SRVDC02
Domain Contr.
VLAN 200, AD Site: B
SRVDC03
Domain Contr.
Board Building
High School Building

Summary

- Make sure you spend time on design before you build the solution. With Exchange 2010 and VMware SRM, planning is a must.

Thank you!

For questions after this presentation, email to

ymagalif@promedianj.com