vmware

#VIRT9132

be_ TOMORROW

VIRT9132

Virtualized Extended Distance RAC on Hyperconverged Infrastructure in a Private Cloud - The Perfect Marriage of Availability and On-demand Scalability

Deborah Kearney, CU ATM Services, LLC
Marlin McNeil, Yucca Group

vmworld 2016

The Percentage of Applications in Virtualized Infrastructure has Increased Dramatically over the Last Few Years

(VMware Core Metrics Survey July 2015)

- Microsoft SQL is the most common application running in on-premise virtual infrastructure

Applications in Virtualized Infrastructure
(Select all that apply)

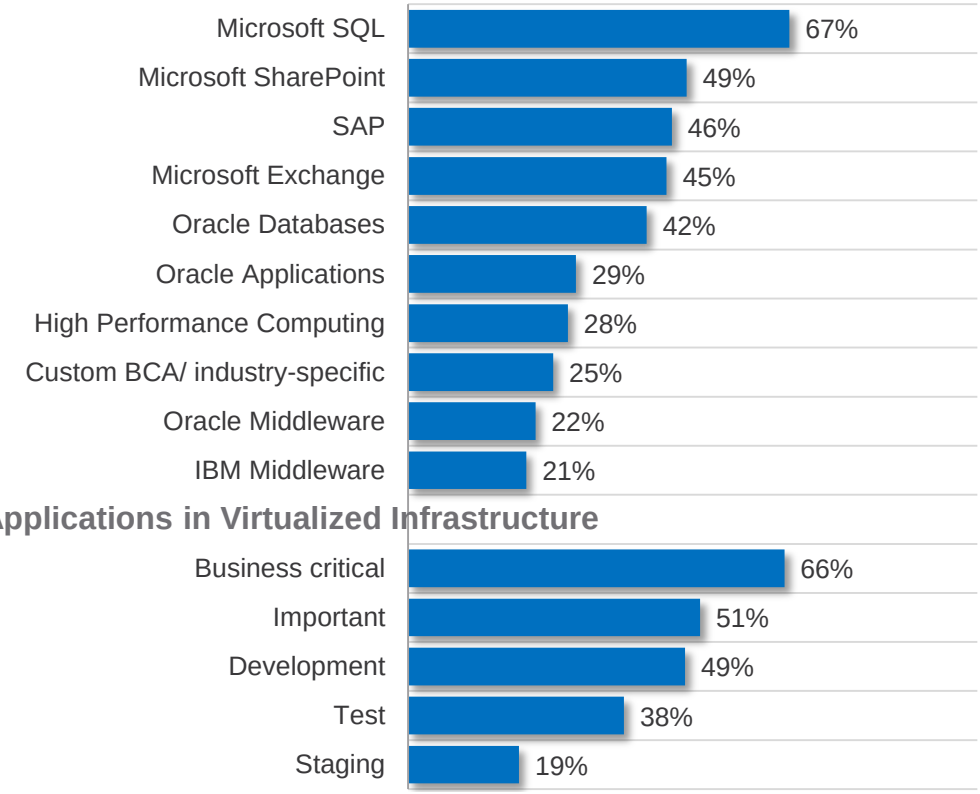

Level of Criticality of Applications in Virtualized Infrastructure
(Select all that apply)

N = 1603

	Region				Company Size		
	NA	EU	dAP	BRIC	SMB	COMM	ENT
Microsoft SQL	**57%**	**73%**	70%	**74%**	68%	71%	64%
Microsoft SharePoint	47%	51%	**39%**	**56%**	**43%**	51%	**54%**
SAP	**41%**	43%	46%	**61%**	**36%**	46%	**57%**
Microsoft Exchange	45%	**54%**	**37%**	41%	43%	49%	46%
Oracle Databases	**34%**	38%	**59%**	**51%**	**37%**	39%	**48%**
Oracle Applications	26%	27%	32%	**37%**	**24%**	**34%**	33%
High Performance Computing	25%	30%	23%	**35%**	**16%**	30%	**39%**
Custom BCA/ industry-specific	**29%**	**16%**	**31%**	27%	22%	22%	**30%**
Oracle Middleware	**15%**	23%	**30%**	**28%**	**19%**	24%	25%
IBM Middleware	**15%**	22%	22%	**30%**	**17%**	21%	**25%**
Business critical	**71%**	62%	62%	64%	65%	64%	68%
Important	48%	54%	49%	55%	50%	51%	53%
Development	51%	45%	49%	49%	**44%**	49%	**53%**
Test	36%	35%	39%	**46%**	37%	40%	37%
Staging	20%	**15%**	20%	**26%**	**15%**	17%	**25%**
	600	450	230	323	653	346	604

> Total
< Total

Virtualizing Applications Sessions and Offerings

- Spotlight Session – Travis Credit Union – 100% Virtual w/Oracle
- 30 Breakout Sessions with 2 Panels & 3 Quick Talks
- 12 BCA Experts for Meet-The-Experts individual 15 minute Appts
- 5 Group Discussions
- 2 Saturday Day Long Bootcamps

Sign up for the Independent Oracle User Group (IOUG) VMware Special Interest Group (SIG)

www.ioug.org/vmware

Prominent Speakers

- **Dave Welch** – Oracle Licensing
- **Allan Hirt** – SQL Server HA & DR with SRM on vSphere
- **Denny Cherry** – SQL Server in the Clouds
- **David Klee & Tom LaRock** – SQL on vSphere Performance
- **Mike Corey** – Doing IT Right
- **Rommel Garcia** – Big Data
- **Todd Muirhead, Reza Taheri, Mark Achtemichuk**

2016 Survey of IOUG Members on Virtualization Adoption and Use

- The perception of increased Licensing and Support costs are the top challenges for organizations seeking to expand the number of Oracle databases and applications being deployed on vSphere
- Yet the Survey found that cost reduction is the top benefit reported amongst virtualization users along with greater consolidation and agility
- VMware vSphere is the number one virtualization solution in use by Oracle shops – representing 70% of IOUG members in the survey
- VMware vSphere for Oracle is expanding:
  - 32% plan to expand virtualization within the 6 months
  - 47% plan to expand virtualization within the 12 months
  - 69% plan to expand virtualization within the 24 months

RDBMS Books from VMware Press

http://www.pearsonitcertification.com/store/virtualizing-oracle-databases-on-vsphere-9780133570182

http://www.pearsonitcertification.com/store/virtualizing-sql-server-with-vmware-doing-it-right-9780321927750

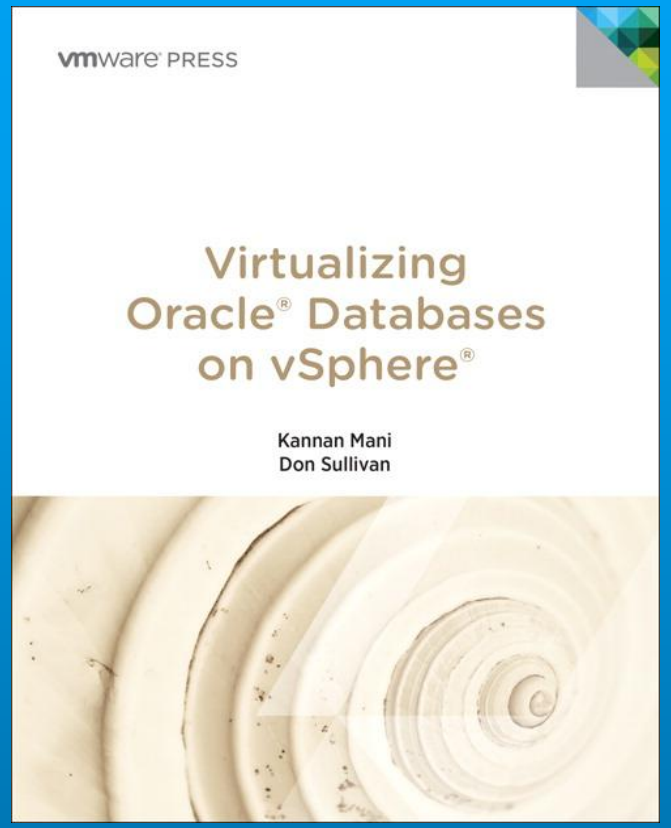

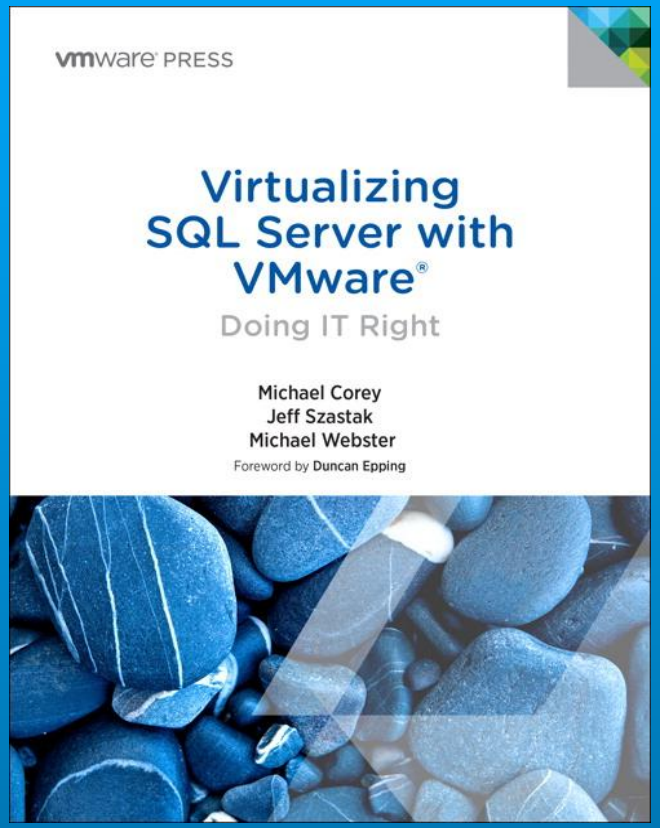

Oracle on vSphere
Book signing @ 1PM
Tuesday Aug 30

vmwarepress.com

Virtualized Extended Distance RAC on Hyperconverged Infrastructure in a Private Cloud

The perfect marriage of availability and on demand scalability

Deborah Kearney, CU ATM Services, LLC

Marlin McNeil, Yucca Group

Introduction

Deborah Kearney

Deborah Kearney is the Chief Information Officer of CU Anytime, LLC, and its subsidiary CU ATM Services, LLC, a New Mexico based Credit Union Service Organization (CUSO). Her role as CIO includes leading the company's technology strategy, management of technology infrastructure, business continuity planning, data security and compliance. She also oversees the company's 3rd party technical support center for remote deposit capture (RDC) image customers. Deborah has 15+ years of experience managing technology in the credit union industry.

www.cuanytime.org

Marlin McNeil

- **20+ years Oracle DBA/Architect**, Oracle RAC, Experienced in Extended Distance Clusters, Data Center Design, PCI Compliance
- **Network Engineer** (Sandia National Laboratories contract) (1991 – 1996)
- **Principal Oracle DBA/Oracle Architect/Technical Services Director** - Justice Systems, Inc. (1996 – 2011)
- **Co-Founder, CTO, Senior Solution Architect - Oracle** – The Yucca Group (2011-Present)
- **Speaker** @ IOUG Collaborate, VMWorld, VMware Partner Exchange, New Mexico VMUG
- **Technical Editor** for "Virtualizing Oracle on vSphere"
- Blog: http://yuccagroup.com/index.php/yg-blog

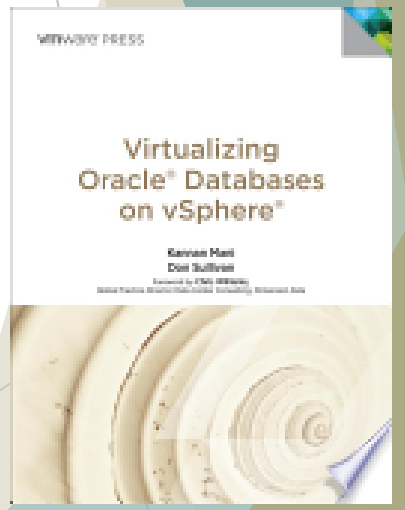

Agenda

- Introduction
- Business Environment
- Requirements
- Hyperconverged Architecture
- SimpliVity OmniCube Infrastructure
- OmniCube and Oracle Considerations
- Existing Resources
- Design Considerations
- Architecture – CU Anytime
- Other Enterprise Features
- RAC on SimpliVity Storage – Issues
- Oracle Licensing
- Performance
- Summary
- Q&A

Business Environment

Business Environment

CU Anytime is a Credit Union Service Organization (CUSO) managing a fleet of 215+ ATMs for 24 New Mexico and El Paso, Texas based Credit Unions.

- Business Goal: Scaling to 5000+ ATMs
- High availability required – Goal of 98% uptime
- Reporting of performance and uptime metrics critical in real time
- Real time incident reporting and dispatching
- Remote visibility and inventory
- Cash allocation forecasting

Requirements

Requirements

Installation of Fault Monitoring and Cash Allocation software platforms requiring the necessity of high availability Windows Server and Oracle based applications.

- Starting at 215+, scalable to 5,000+ ATMs
- For all applications
  - Recovery Point Objective (RPO) = 0
  - Recovery Time Objective (RTO) = 30 sec
- Load requirements
  - Base – ATM Transaction, Fault and Cash Load information
  - Scaled - By a factor of 25

Hyperconverged Architecture

Hyperconverged Infrastructure (HCI)

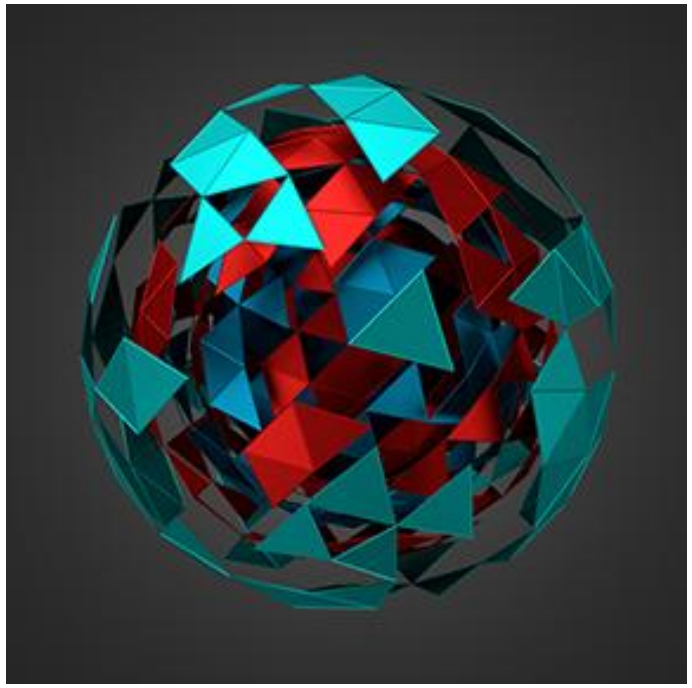

A convergence of solutions and reduction of complexity

- Integrates computing, storage, networking, virtualization resources and other technologies in a hardware box supported by a single vendor

Unified management

- Converged management interface allows for configuration of all included systems from single management console

SimpliVity OmniCube Infrastructure

SimpliVity's OmniCube Infrastructure

- Data Efficiency – Global inline deduplication, compression and data optimization
- Built-In Data Protection – Virtual Machine (VM)-centric local and remote backups
- VMWare's vCenter Management Console – SimpliVity resources and features are centrally provisioned from vCenter and the vSphere management console
- Supports Virtual Desktop Infrastructure (VDI)

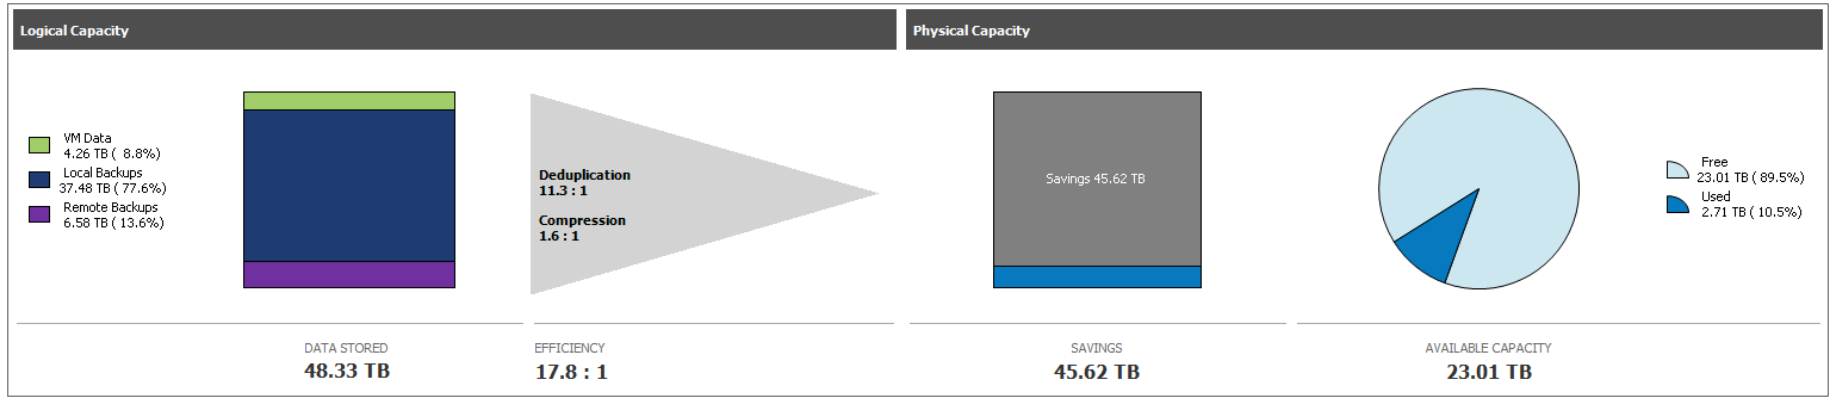

SimpliVity's OmniCube Infrastructure

OmniCube CN-3400

- CPU Config: 12 Cores per socket, two sockets represent 24 cores
- Memory: 386G
- Usable Disk Space: 20TB raw, conservatively usable = 13TB
- Logical Disk Space: 13TB * 3 = 39TB

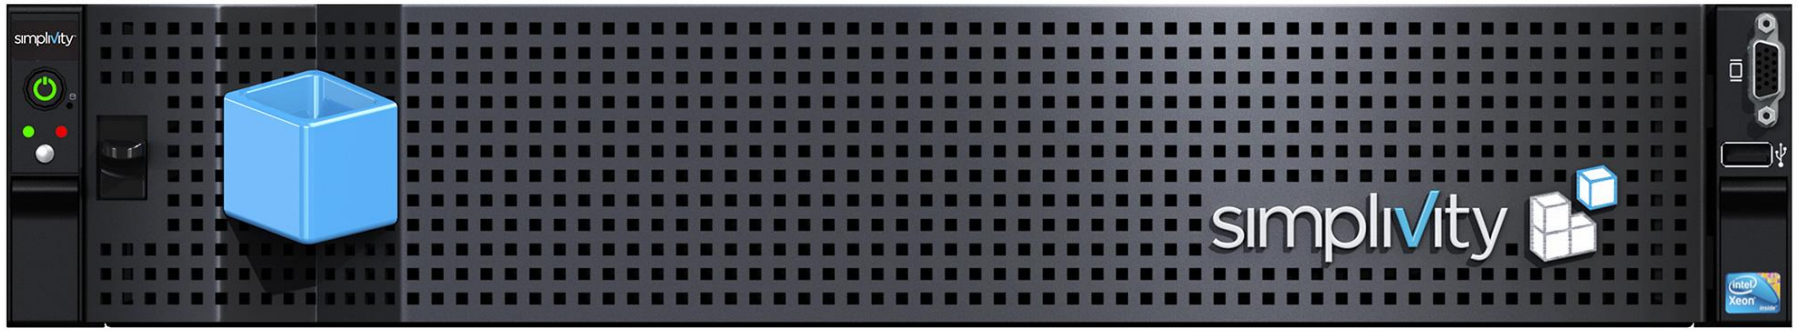

Leveraging OmniCube Features

OmniCube Shared Storage and Oracle

- SimpliVity Hardware Accelerator
  - Fast IO, cache efficiencies improve data access speed
  - Deduplication and compression
  - Backup capabilities
- Network File System (NFS) storage accessible from compute vSphere hosts
  - Oracle Licensing Considerations

Existing Resources

Existing Infrastructure Resources

- Two sites located 6 miles apart
- 50 Mbps fiber between sites
  - 1ms round trip latency
- Two (2) SimpliVity OmniCube CN-3400 at Site 1
- One (1) SimpliVity OmniCube CN-3400 at Site 2
- Confederated OmniCube storage
  - 24TB at Site 1
  - 12TB at Site 2
  - Mounted over NFS
  - vSphere 6.0 Update 1
- Two (2) repurposed servers available after virtualization (one at each site)
  - Server specs
    - 4 Cores, 3 GHz, 40GB DDR4 2133, Four (4) 1 Gigabit Interfaces
    - 6 Cores, 2.4 Ghz, 40GB DDR3 1333, Four (4) 1 Gigabit Interfaces
  - vSphere 6.0 Update 1

Design Considerations

Traditional Cluster Deployments

- Connected computers - function as single unit
- High availability (HA)
- Scale
  - Horizontally by adding nodes to cluster
  - Vertically by swapping out smaller nodes with beefier nodes
- Typically in single data center
  - Historically due to lack of WAN bandwidth
- Cluster Requirements
  - Shared storage required
  - 1 Gbps / 10 Gbps (Jumbo Frames recommended)
    - vSphere vMotion LAN
    - Shared Storage access

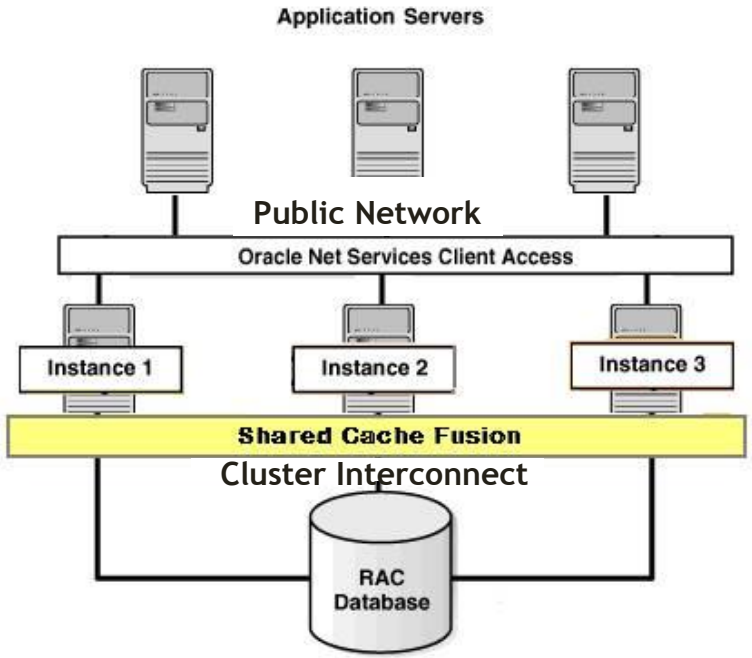

Extended Distance Clusters

- Provides High Availability (HA)
- Scales horizontally and vertically
- Stretched across multiple data centers
  - Storage and network stretched across sites
  - Universally accessible from all sites
- Cluster requirements
  - Still the same
    - Replicated Shared Storage
    - WAN network bandwidth for
      - Oracle RAC Private Interconnect, vSphere vMotion LAN, Shared Storage access
- Latency
    - Oracle RAC - 5 ms
    - vSphere vMSC – 5-10 ms
    - Storage (most vendors) – 2-5 ms
- Layer 2 Adjacency

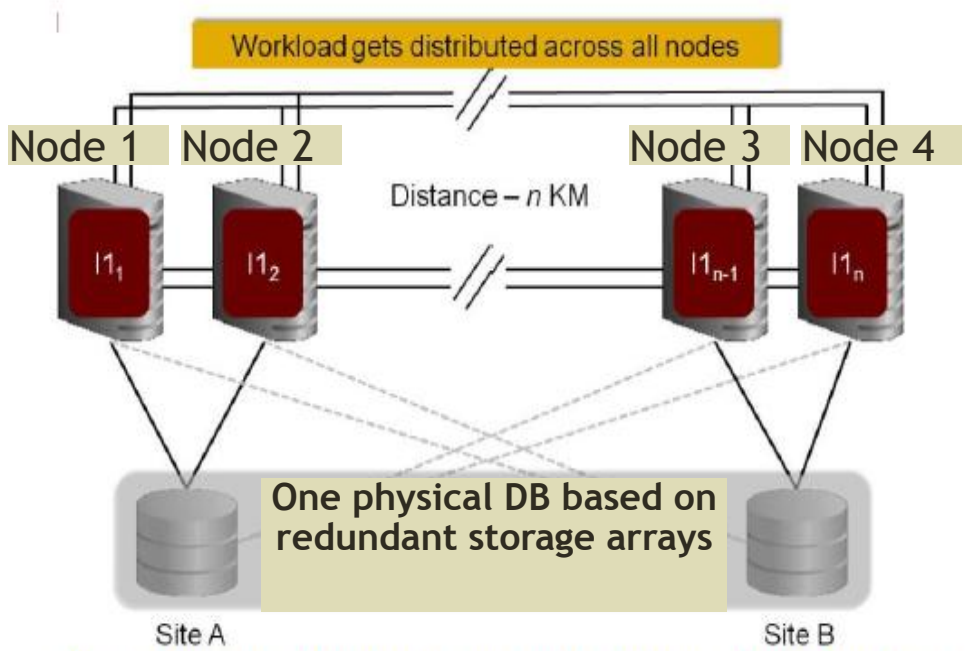

Figure 2: Distributed Workload in an Extended Distance Oracle RAC Cluster

Oracle RAC and Oracle RAC One Node on Extended Distance (Stretched) Clusters

Extended Distance Clusters – Use Cases

High availability across locations

- Protection from local disaster with minimal service disruption
- Disaster Avoidance instead of Disaster Recovery

Workload balancing across sites

- Optimal use of resources in both sites
- Migrate workload from one site to another for load balancing

Site maintenance without service disruption

- Orchestrated site evacuation
- Universal Protection for Applications
- Application independent, no need for additional tools

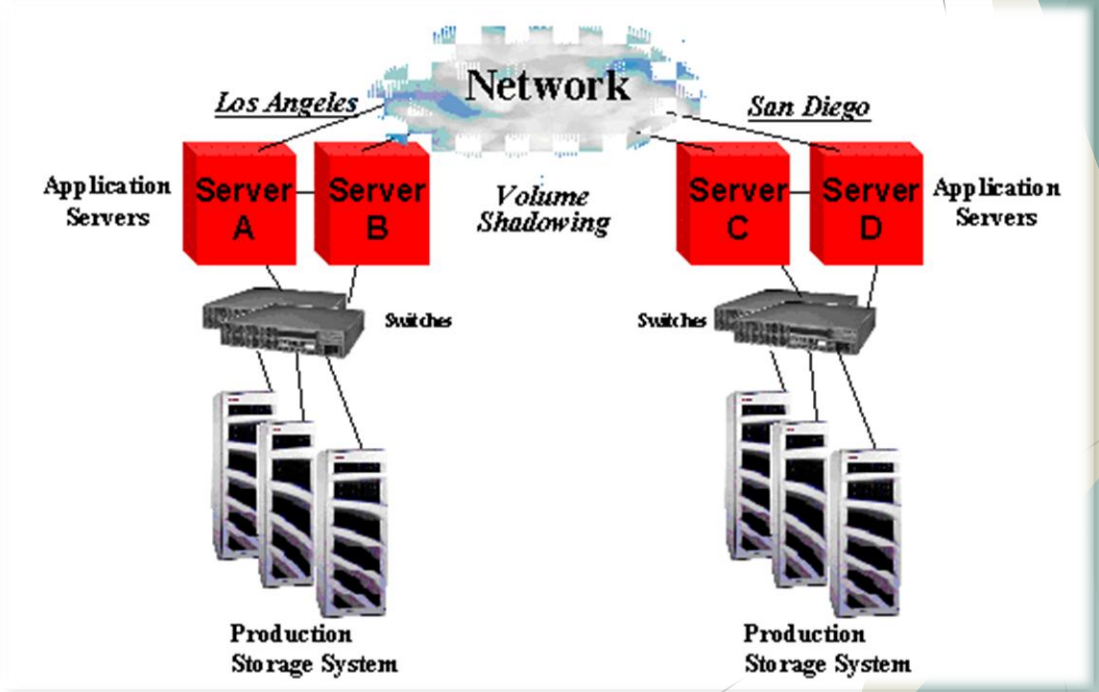

http://h71000.www7.hp.com/openvms/journal/v2/articles/cluster.html

Extended Distance Clusters – Misconceptions

This is disaster avoidance - not disaster recovery

- HA is automatic, manual intervention not required
- No run book – does not have plan based actions & reporting

Clustering is across sites, not components

- Third-party witness site required

Not necessarily the ‘cheaper’ solution

- Costs related to storage subsystem and complexity
- Higher network bandwidth requirements than a DR solution

Extended Distance – Characteristics of Data Communication

- Speed of light – Challenges in physics
  - 300,000 KM/sec in vacuum
  - 200,000 KM/sec in Fiber Optic (200 KM/ms or 124 miles/ms)
- Latency characteristics in Extended Distance Clusters
  - Round trip
  - Cable distance
  - Includes router hops
    - Store and forward
  - Rule of thumb
    - Latency less than 5ms is stable
- Example: Latency considerations for Oracle RAC
  - Stable Up to 50 KM
  - Possible Up to 100 KM – Test the Performance first
  - Refer “Oracle RAC and Oracle RAC One Node on Extended Distance (Stretched) Clusters”

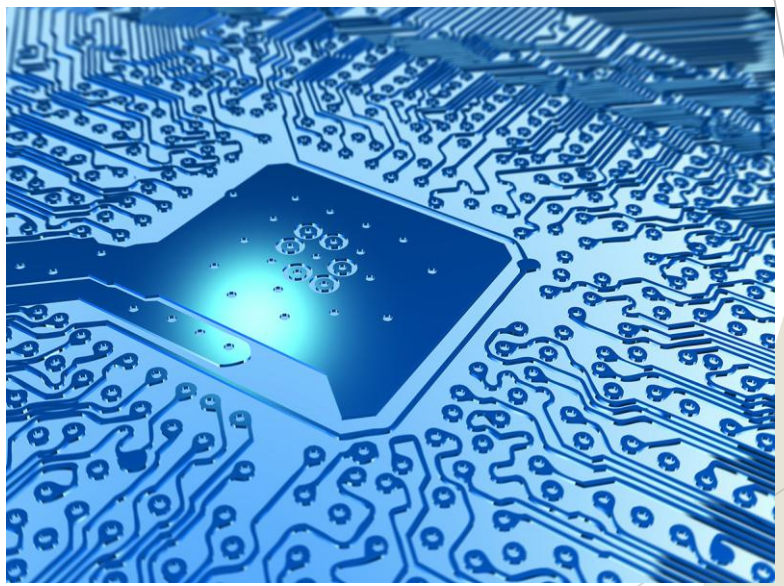

Stretched Oracle RAC Using Automatic Storage Management (ASM) Based Mirroring

- A third voting file on a standard NFS or iSCSI mounted device at a third site
- Recommend putting the 3rd voting file on a dedicated server/SAN which belongs to a production environment

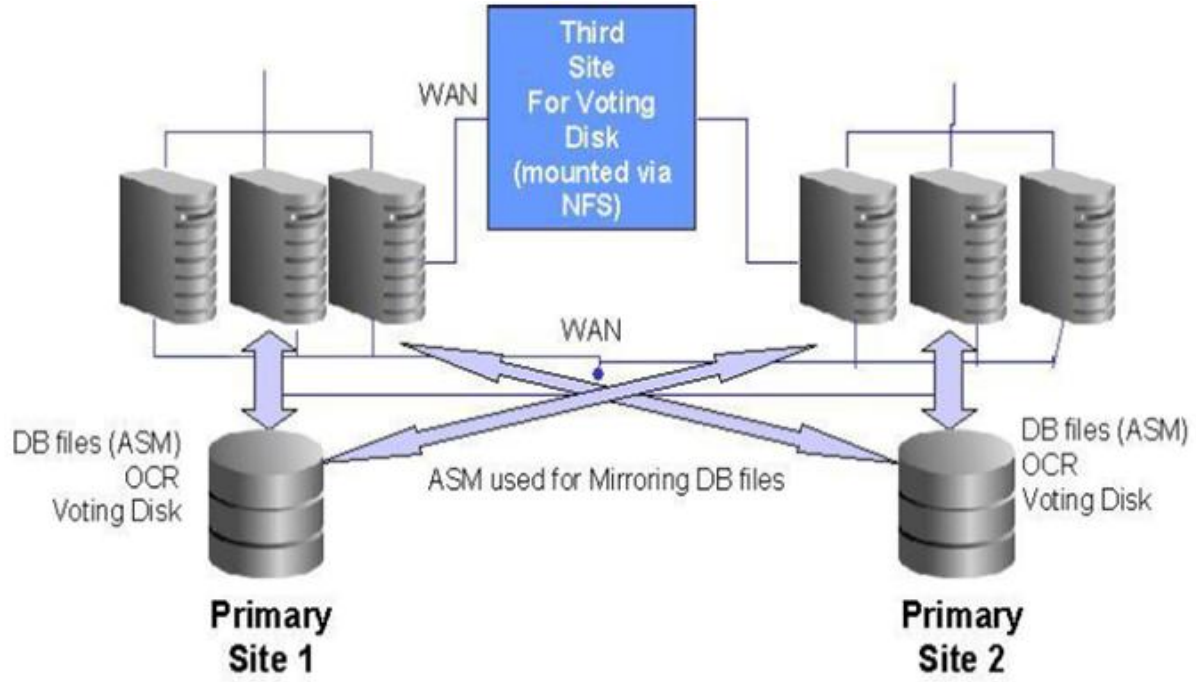

Figure 2: Extended RAC environment with standard NFS Voting file in third site

Oracle Clusterware 11g Release 2 (11.2) – Using standard NFS to support a third voting file for extended cluster configurations

31

Architecture – CU Anytime

Extended Oracle RAC

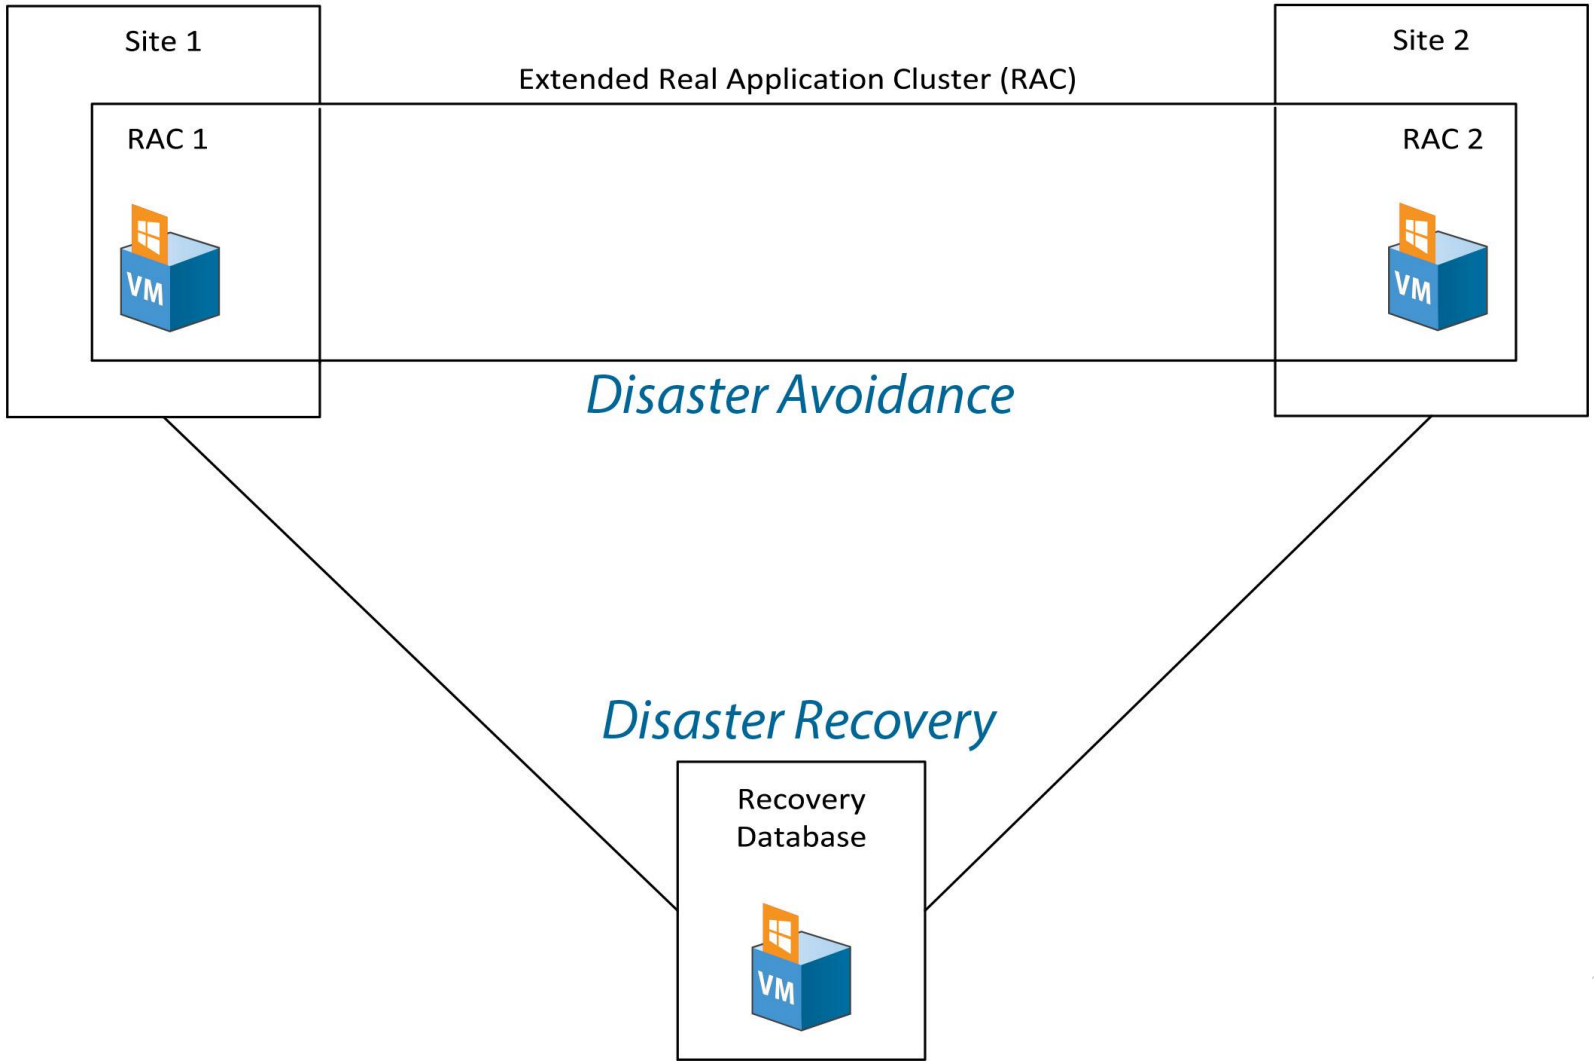

Oracle Environment

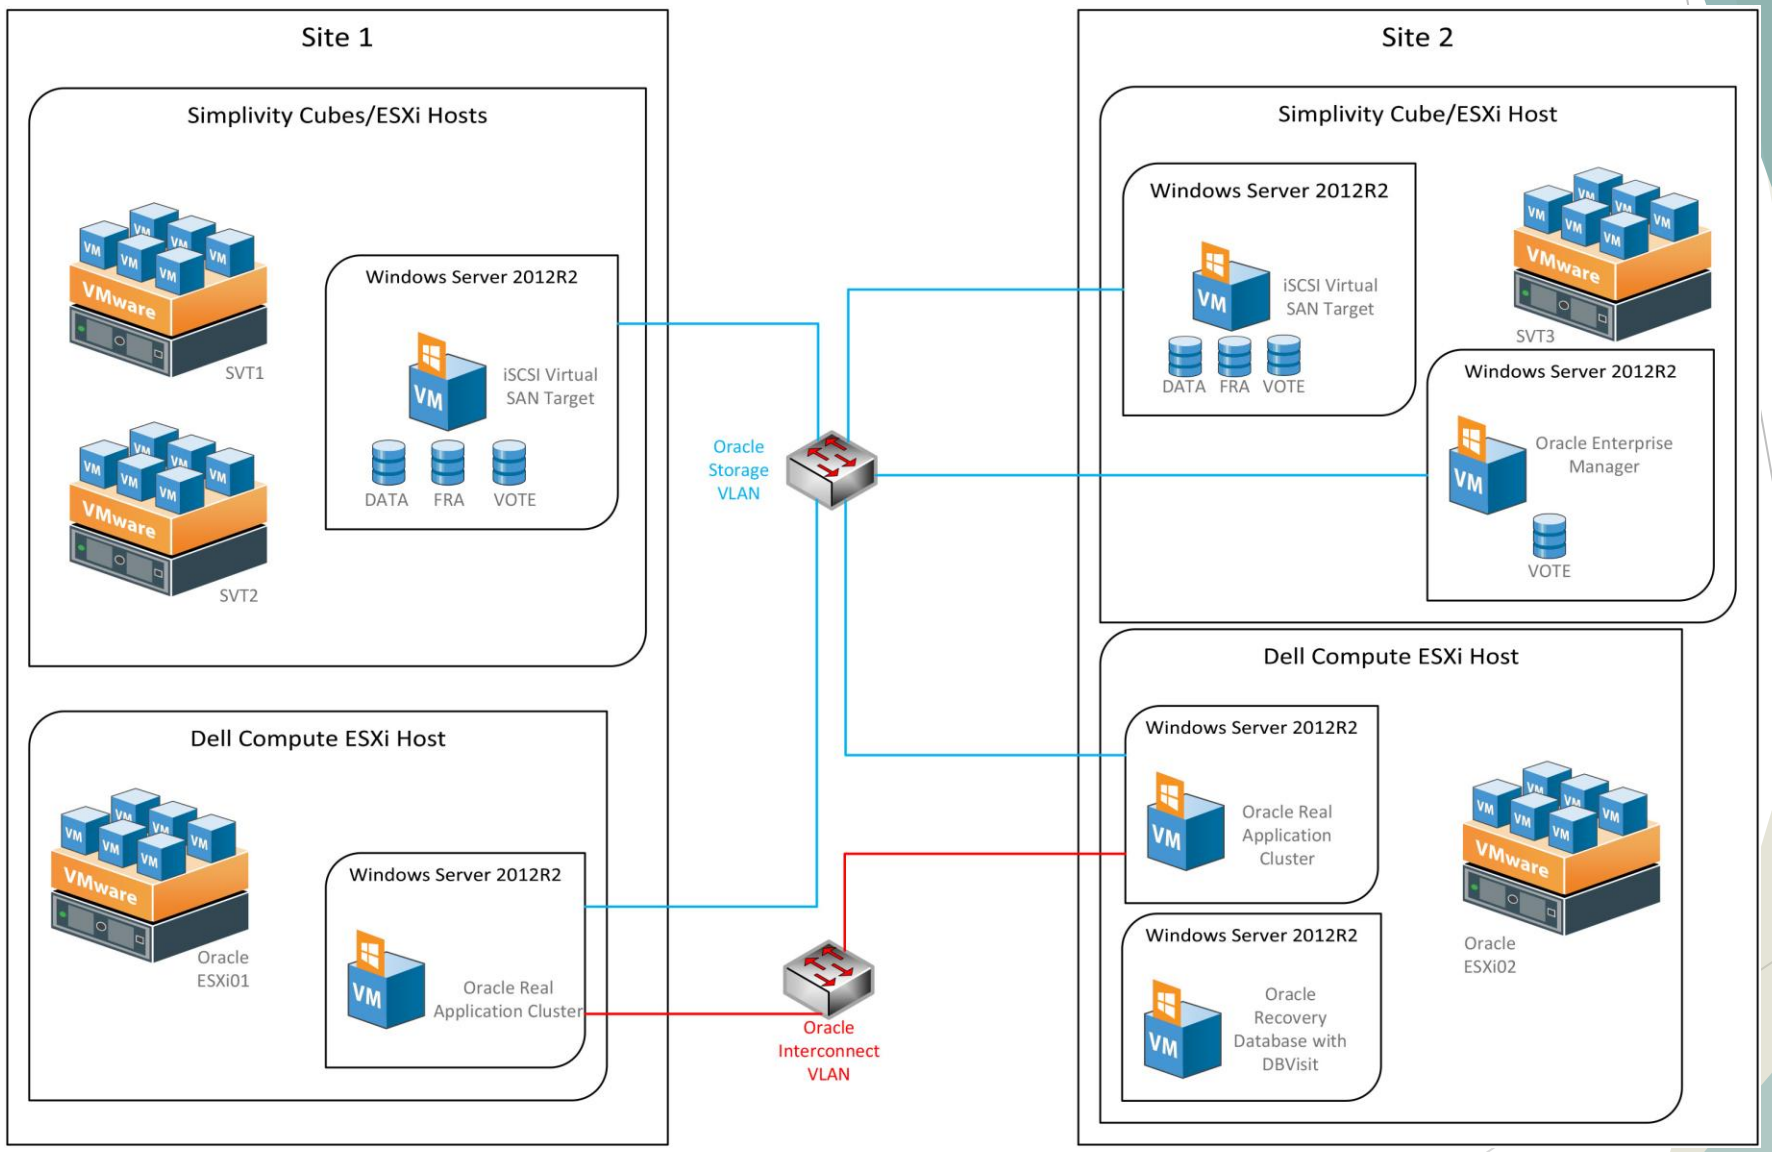

Site Architecture

- Clustered OmniCubes at Site 1
- Single OmniCube at Site 2
- Each site sends OmniCube backup to other site
- Virtualized extended distance Oracle RAC between sites
  - Windows Server 2012 R2 – RAC nodes, iSCSI targets
  - Oracle 12C SE2
  - Single RAC node at each site
- Layer 2 adjacency using Cisco L2TP
- Extended VLAN for public, interconnect networks
- iSCSI routed between sites
  - Reduces latency between sites
- Storage replication for RAC accomplished with Normal redundancy ASM Disk Groups
- Disaster Recovery Site using DBVisit Standby
- Performance Tuning and Monitoring using SolarWinds Data Performance Analyzer (DPA)

Other Enterprise Features

Disaster Recovery – DBVisit Standby

- DBVisit Standby can work with Oracle SE2
- Same functionality as Oracle Active Data Guard

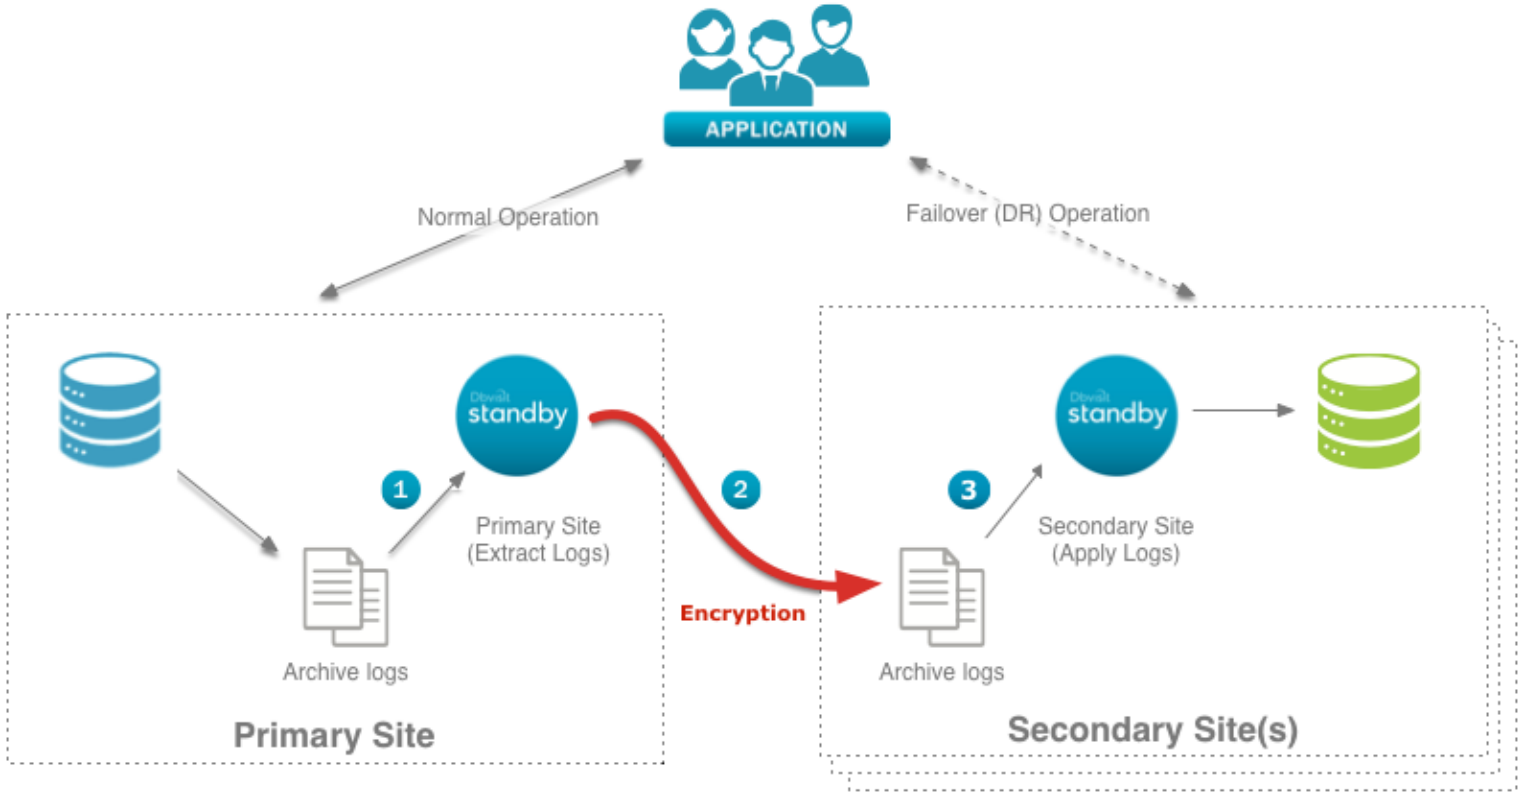

Tuning and Monitoring – SolarWinds DPA

- SolarWinds Database Performance Analyzer (DPA)
  - Performance tuning
  - Monitoring
- Same functionality as Oracle Diagnostics and Tuning pack

RAC on SimpliVity Storage - Issues

RAC on SimpliVity Storage - Issues

SimpliVity storage does not natively support RAC shared storage

THE SOLUTION: Virtualize storage – Windows Server 2012 iSCSI target

- Masks incompatibility

HOW: In-guest iSCSI LUNs in ASM disk groups

- ASM allows storage replication between sites for extended distance RAC

RAC on SimpliVity Storage - Issues

Problems you don't have to solve yourself…

- Set the Local Adapter and Source IP Address in the Target portals (under the Discovery
  tab in the iSCSI initiator), using the Advanced button
- Set SAN policy on RAC nodes to OnlineAll
- Needs to be on MBR partitions (critical but inconspicuously documented)
  Grid Infrastructure Installation Guide Section 6.6.1
  Create Disc Partitions for Use With Oracle ASM
- Silent OEM agent install on Windows hosts with command line arguments

RAC Interconnect - Issues

- 50mb Interconnect
  - Not sufficient to carry significant cache fusion traffic
- Used single database instance protected by RAC
- Used Oracle services
  - Local instance – preferred
  - Remote instance – available
- Result - No cache fusion traffic across interconnect

Oracle Licensing

Oracle Licensing

- Oracle 12c (12.1.0.2) Database SE2 RAC
- RAC node VMs reside on repurposed hardware
- All CPUs on repurposed hardware are licensed for Oracle
  - Single server at each site
  - Each server has a single populated socket
- RAC VMs confined to compute server (not SimpliVity OmniCube) using node affinity
- NFS volumes containing Oracle binaries mounted to “compute” server and one SimpliVity OmniCube per site (for standby server)
- SimpliVity OmniCube serves as standby for Oracle “10 day rule”

Performance

Performance – vMotion Live Migration

17232 tpm

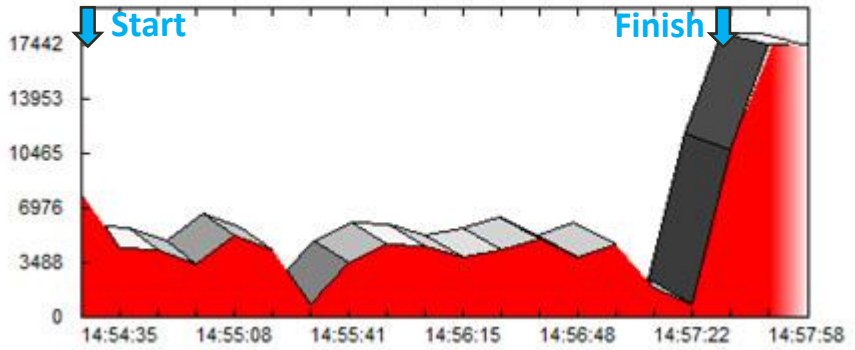

Graph 1 - Compute Node to OmniCube

22050 tpm

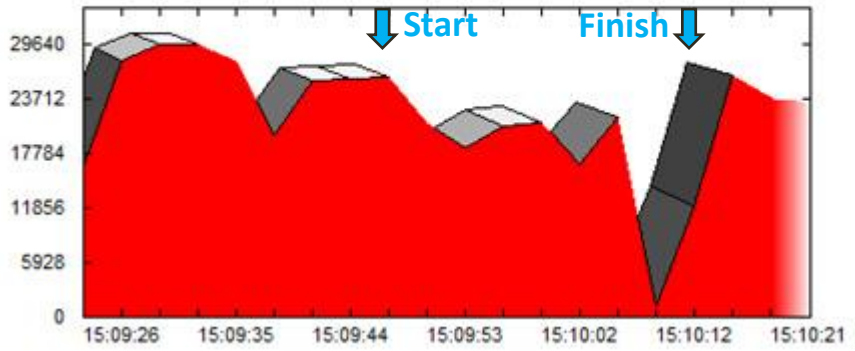

Graph 2 - OmniCube1 to OmniCube2

Performance – RAC Interconnect

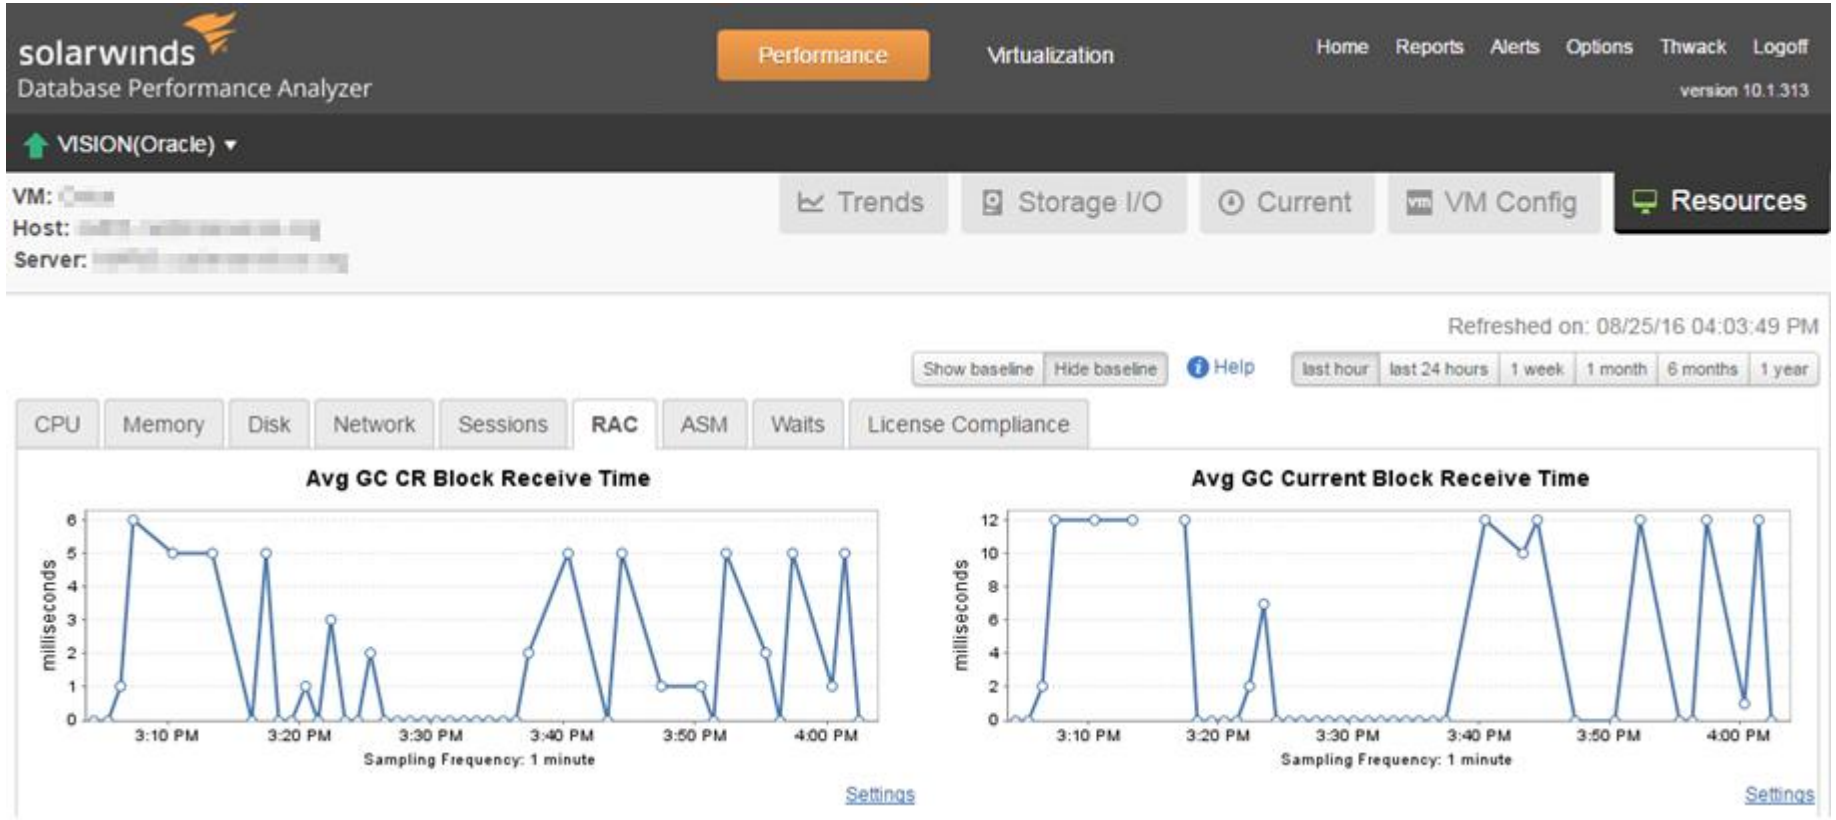

Performance – Disk IO

DBMS_RESOURCE_MANAGER.calibrate_io

- Max IOPS = 29256
- Max MBPS = 634
- Latency = 7

dba_rsrc_io_calibrate

Query Result

SQL | All Rows Fetched: 1 in 0.004 seconds

	START_TIME	END_TIME	MAX_IOPS	MAX_MBPS	MAX_PMBPS	LATENCY	DISKS
1	25-AUG-016 16:38:32	25-AUG-016 16:51:31	29256	634	542	7	1

Conclusion

Conclusion

- Business critical applications can be successfully virtualized using hyperconverged infrastructure
- Storage virtualization can mitigate issues with underlying physical storage
- On hyperconverged infrastructure, applications can scale without large capital expenditure
- Oracle SE2 licensing can be leveraged for a SMB environment
- Oracle enterprise style features leveraged through 3rd party add-ons

Q&A

vmware

THANK YOU

*be*_TOMORROW

vmworld 2016

vmware

FILL OUT YOUR SURVEY

Take a survey and enter a drawing for a VMware company store gift card

vmworld 2016

vmware

#VIRT9132

be_ TOMORROW

VIRT9132

Virtualized Extended Distance RAC on Hyperconverged Infrastructure in a Private Cloud - The Perfect Marriage of Availability and On-demand Scalability

Deborah Kearney, CU ATM Services, LLC
Marlin McNeil, Yucca Group

vmworld 2016

References

- VMWare KB 2031038 - VMware vSphere 5.x and 6.0 support with NetApp MetroCluster
  https://kb.vmware.com/kb/2031038
- VMWare KB 2000767 - Unable to initialize a second virtual disk within Windows 2008 or Windows 2008 R2
  https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2000767
- Oracle - High Availability Architectures and Solutions
  http://docs.oracle.com/cd/B28359_01/server.111/b28281/architectures.htm
- Oracle - Configuring Oracle Database with Oracle RAC
  http://docs.oracle.com/cd/E11882_01/server.112/e10803/config_rac.htm#HABPT4874
- Oracle - Configuring Shared Storage for Oracle ASM
  https://docs.oracle.com/database/121/CWWIN/storage.htm#CWWIN316
- IBM Redbooks - SAN and SVC Stretched Cluster and VMWare Solution Implementation
  http://www.redbooks.ibm.com/redbooks/pdfs/sg248072.pdf
- IBM Redbooks - Spectrum Virtualize and SAN Volume Controller Enhanced Stretched Cluster with VMWare
  http://www.redbooks.ibm.com/redbooks/pdfs/sg248211.pdf
- Dbvisit Standby
  http://www.dbvisit.com/products/dbvisit_standby_database_for_oracle_disaster_recovery/
- Microsoft Technet - iscsi initiator target not reconnecting on reboot
  https://social.technet.microsoft.com/Forums/windowsserver/en-US/4b2420d6-0f28-4d12-928d-3920896f582d/iscsi-initiator-target-not-reconnecting-on-reboot?forum=winserverfiles