



# PRE-CONFERENCE SEMINAR

HOSTED BY



## **Sneak Peek:** Disaster Proof: Using OpenVMS Clusters for Disaster Tolerance

Keith Parris  
Systems/Software Engineer  
May 6, 2009



# Overview

- Trends affecting Disaster Tolerance / DR
- Common questions
- Technology options, and which are most common in practice
- Crucial design choices

# Trends affecting Disaster Tolerance / DR

- World-wide trends
- Computer Industry trends
- Technology trends
- DR/DT configuration trends
- OpenVMS technology directions



# World-wide trends

- Increase in disasters
- Economic downturn

# Computer industry trends

- Cost-cutting
- Hardware faster, cheaper, smaller, denser
- Network links faster, cheaper

# DR/DT configuration trends

- Longer inter-site distances for better protection
- Business pressures for shorter distances for better performance
- IP network focus; increasing pressure not to bridge LANs between sites
- Harder to get dark fiber; easier to get lambdas (DWDM channels)

# Technology trends

- Disk drive capacity going up; cost going down
- Interconnects:
  - High end: Ethernet + Fibre Channel, plus FCIP
  - Low end: Ethernet + SATA/SAS/SCSI or iSCSI
  - Looking at Converged Enhanced Ethernet and FCoE
- Blades
- Virtualization

# OpenVMS technology directions

- IP routing as network sweet spot, not bridging [thus OpenVMS will support IP as a Cluster Interconnect]
- Desire for redundancy after site failure [thus support for more-than-3-member shadowsets]

# Common questions



- Questions people often ask when considering a disaster-tolerant OpenVMS Cluster configuration

# Common questions when considering a disaster-tolerant OpenVMS configuration

- My current cluster uses a quorum disk. Can I use Continuous Access to mirror the quorum disk between two sites?
  - What's a quorum node?

# Common questions when considering a disaster-tolerant OpenVMS configuration

- Do I need Disaster Tolerance or is Disaster Recovery good enough?

# Common questions when considering a disaster-tolerant OpenVMS configuration

- Should I have a multi-site cluster or just use data replication (e.g. Shadowing) between sites (“Remote Vaulting”)

# Common questions when considering a disaster-tolerant OpenVMS configuration

- Should I use Continuous Access or Host-Based Volume Shadowing?

# Common questions when considering a disaster-tolerant OpenVMS configuration

- How far apart can my sites be?
  - Can I use my existing datacenter sites in cities X and Y?

# Common questions when considering a disaster-tolerant OpenVMS configuration

- Do I need three sites or just two?

# Common questions when considering a disaster-tolerant OpenVMS configuration

- How far apart should my sites be for safety?

# Common questions when considering a disaster-tolerant OpenVMS configuration

- How much adverse impact will there be on application performance with sites a distance of X apart?

# Common questions when considering a disaster-tolerant OpenVMS configuration

- What quorum scheme should I use? (balanced, unbalanced, quorum disk, quorum node?)

# Common questions when considering a disaster-tolerant OpenVMS configuration

- How much bandwidth do I need between sites?
  - How do I calculate or measure or somehow determine how much bandwidth I need?

# Common questions when considering a disaster-tolerant OpenVMS configuration

- What settings should I use for various SYSGEN parameters?:
  - RECNXINTERVAL
  - TIMVCFAIL, PE4
  - SHADOW\_MBR\_TMO
  - LOCKDIRWT, LOCKRMWT
  - VOTES, EXPECTED\_VOTES, DISK\_QUORUM, QDSKVOTES, QDSKINTERVAL
  - SHADOW\_SYS\_DISK, SHADOW\_SYS\_UNIT
  - MSCP\_CREDITS, MSCP\_BUFFER
  - CLUSTER\_CREDITS

# Common questions when considering a disaster-tolerant OpenVMS configuration

- With  $X$  terabytes of data and a link with bandwidth of  $Y$ , how long will it take for a shadowing Full-Copy operation? How about for a Mini-Copy?

# Common questions when considering a disaster-tolerant OpenVMS configuration

- What happens when I lose the inter-site LAN link?
- How about the inter-site SAN link?
- What if I lose both at the same time?

# Common questions when considering a disaster-tolerant OpenVMS configuration

- I have a long-distance OpenVMS cluster with 20 miles between sites. To help guide Volume Shadowing, I'll be using different SHADOW\_SITE\_ID SYSGEN parameters for each site (or different Read\_Cost values) so all the reads will be to disks local to each site. Is that right?

# Common questions when considering a disaster-tolerant OpenVMS configuration

- I have FCIP SAN Extension boxes in place so the systems in the cluster at both sites can see all the disks on both SANs. Since all my systems have direct access to all the disks, why in the world would I want to enable MSCP Serving?

# Common questions when considering a disaster-tolerant OpenVMS configuration

- Some of my disks contain data which is rarely changed. Should I plan to restore those disks from backups in the event of a disaster, or just shadow all my disks for the sake of consistency?

# Common questions when considering a disaster-tolerant OpenVMS configuration

- Should I shadow the system disk across sites?

# Common questions when considering a disaster-tolerant OpenVMS configuration

- I'd like to have two shadowset members per site (and I have 2, or 3 sites). How can I do that with a limit of 3 shadowset members?

# Common questions when considering a disaster-tolerant OpenVMS configuration

- My network folks refuse to allow any bridging between sites. How can I set up a multi-site OpenVMS disaster-tolerant cluster?

# Technology options

- Various technology alternatives available, and what is typical in configurations seen in the real world



# Technology options, and which are most commonly found in practice

- Disk replication by disk controllers (i.e. Continuous Access or SRDF), or
- Host-based Volume Shadowing

# Technology options, and which are most commonly found in practice

- Disk replication by disk controllers (i.e. Continuous Access or SRDF), or
- **Host-based Volume Shadowing**

# Technology options, and which are most commonly found in practice

- Data replication only (remote vaulting), or
- Multi-site cluster


# Technology options, and which are most commonly found in practice

- Data replication only (remote vaulting), or
- **Multi-site cluster**

# Technology options, and which are most commonly found in practice

- Two-site cluster configuration, or
- 3-site cluster configurations

# Technology options, and which are most commonly found in practice



# Technology options, and which are most commonly found in practice

- **Two-site cluster configurations (85%)**
- 3-site cluster configurations (15%)
- *IP as a Cluster Interconnect may help raise 3-site numbers significantly, as it will make having a quorum node at a 3<sup>rd</sup> site much easier*

# Technology options, and which are most commonly found in practice

- Inter-site links:
  - DS-3 [E3 in Europe], or
  - OC-3, or
  - OC-12, or
  - OC-48, or
  - Dark fiber, or
  - Lambdas (individual channels over DWDM)

# Technology options, and which are most commonly found in practice

- Inter-site links:
  - DS-3 [E3 in Europe]
  - **OC-3** (sweet spot for WANs)
  - OC-12
  - OC-48
  - **Dark fiber** (sweet spot for sites sharing same campus)
  - **Lambdas** (sweet spot for MANs)

# Technology options, and which are most commonly found in practice

- Local SCSI or SAS disks, or
- Fibre Channel SAN-based storage

# Technology options, and which are most commonly found in practice

- Local SCSI or SAS disks, or
- **Fibre Channel SAN-based storage**

# Technology options, and which are most commonly found in practice

- Servers, or
- Blades, or
- Virtual Machines

# Technology options, and which are most commonly found in practice

- **Servers** (sweet spot for **existing** clusters)
- **Blades** (sweet spot for **new** installations)
- Virtual Machines (future possibilities for DR/DT)

# Crucial choices

- Some of the critical decisions one must make in planning, designing, and implementing a disaster-tolerant OpenVMS configuration



# Critical choices when planning an OpenVMS disaster-tolerant cluster

- What are the business requirements in terms of RPO and RTO metrics?

# Critical choices when planning an OpenVMS disaster-tolerant cluster

- Will I get consulting help or do a roll-your-own solution?

# Critical choices when planning an OpenVMS disaster-tolerant cluster

- Site selection:
  - Locations
  - Distance between sites

# Critical choices when planning an OpenVMS disaster-tolerant cluster

- Multi-site cluster or just remote vaulting?

# Critical choices when planning an OpenVMS disaster-tolerant cluster

- If a cluster: 2 sites or 3?

# Critical choices when planning an OpenVMS disaster-tolerant cluster

- Quorum scheme:
  - Quorum node?
  - Balanced votes between 2 sites? Unbalanced votes?

# Critical choices when planning an OpenVMS disaster-tolerant cluster

- Data replication method:
  - Replication of disk blocks
    - (and then controller-based or host-based?),
  - Database replication
  - Middleware (i.e. RTR)

# Critical choices when planning an OpenVMS disaster-tolerant cluster

- How many redundant copies of data?
  - 1 per site?
  - 2 per site?

# Critical choices when planning an OpenVMS disaster-tolerant cluster

- Remote disk access method:
  - MSCP Serving, or
  - Linking or bridging SANs between sites

# Critical choices when planning an OpenVMS disaster-tolerant cluster

- Inter-site link(s):
  - Type(s)
  - Number

# Critical choices when planning an OpenVMS disaster-tolerant cluster

- Management/Monitoring/Control suite:
  - DTCS package
  - CockpitMgr
  - Roll-your-own solution

# Sample of Additional Topics

- Converting an Existing Cluster to a Disaster-Tolerant Cluster
- Volume Shadowing and Data Protection: “Nightmare” Scenarios, and How to Avoid Them:
  - “Wrong-Way Shadow Copy”
  - “Creeping Doom”
  - “Rolling Disaster”
- Mitigating Long Distances
- Simulation, Testing, and Staff Training
- New Features in OpenVMS for DT Clusters

# Register Today!



**Register for the  
HP Technology Forum &  
Expo 2009 and  
a Pre-Conference Seminar  
that fits  
your needs at**

**[www.hptechnologyforum.com](http://www.hptechnologyforum.com)**

# Thank you for your attention!

## Questions?

You can send your questions to:

[education@connect-community.org](mailto:education@connect-community.org)

