

From Alpha to VSI OpenVMS on Integrity: A case study

**Nic Clews (with thanks to Ian Miller)
Senior Prof. Systems Engineer, OpenVMS Ambassador,
hpUG, dxc.technology**

Background

- What did we do?
- Where does Ian Miller (OpenVMS Ambassador) fit in?
- **Not enough time for a deep dive**
- Design to execution...
- Bumps in the road...
- Questions...

Your guide for this trip: Nic Clews

- “digital” technology experience for just over 40 years (Compaq / HP(E) / VSI)
- Experience across many industries with business critical systems
- Senior Professional Systems Engineer for dxc.technology (formerly CSC merged with HPE services)
- OpenVMS Ambassador, speaker at Bootcamps and hpUG
- Personal interest in high end audio, radio broadcasting, photography and travel, service user representative for Macmillan Cancer support

nclews@dxc.com

Source environment

- Pair of DS25 Alpha systems, OpenVMS 8, EVA storage, SDLT tape
- PRODUCTION located in London, DEVELOPMENT in east England
- User active hours Monday to Saturday morning to early evening
- Daily batch jobs running overnight (including pre and post backup)
 - Database quiesced for save operations
 - Disk to disk, then to tape (SDLT) for speed (manual changes)
- Weekly, monthly and end of year batch processing
- On-call alerting of issues via SMS dial out modem
- Extensive use of logicals to define operating environment

Challenges

- Servers, storage and operating system end-of-support-life
- Local closure of datacentre, server(s) need relocation
- Recent event caused service outage
 - Air conditioning failed
 - Storage array ‘cooked’ and catastrophically failed
 - Remote replication processes failed, unnoticed for months
 - Backup tape en-route to development (off-site) office
 - Lack of documentation on existing environment
 - Second failure (outage) accelerated need for change

Other existing challenges

- Tape changes performed by on site staff
- Weekend full backups meant Fridays were skipped
- Annual tapes (data) need to be kept for 7 years
- Staff retiring
- Ensuring staff felt part of the process, not being replaced by us

Service requirements

- Move from Alpha to Integrity
- Relocate production (and development) servers
- RTO / RPO (Recovery Time Objective / Recovery Point Objective) of 0 (ZERO!)
 - RTO = “zero downtime”
 - RPO = “no data lost”
- Current service non-clustered, no volume shadowing (mirroring), relies on storage system (EVA) data protection

Proposals

- VSI OpenVMS, Integrity rx2800, local storage, LTO backup
 - Development with SDLT to read archive tapes (7 years retained)
 - Only VSI can provide the forward OS support
- Clustered and not clustered
 - Volume shadowed (mirrored) data and development disks
- Backup regime redesign

CONS

- Cost!
- Split site shadowing slowing performance
- Clustering PRODUCTION with DEVELOPMENT is a double edged sword
 - Sharing a UAF (user authorization file), queues, licences...
 - Careful differentiating between the systems
 - Heavy inter-site traffic from development could impact production performance.
- Adequate bandwidth, low latency; critical for performance



PROS

- Real time replication of data to remote site capable of matching 0 RTO / RPO
- Supported hardware and software
- Minimal retraining, retaining proven platform
 - Alpha to Itanium VSI VMS, enhancements, no functional changes
- Remote robotic tape operations
- Zero retraining!
 - Technology familiarization

What did they choose?

In increasing order of cost...

A: Same site standalone

B: Split site standalone

C: Same site clustered

D: Split site clustered

(Split site increases inter site latency, slowing performance)

BONUS QUESTION

VOLUME SHADOWING: Yes or no?

The value of business data

Many people make the mistake of valuing the system in terms of capital, and maintenance.

You really should consider it in terms of the business turnover, the losses due to a failure and loss of data, and the efficiency of the management overhead, time patching and repatching, required downtime, out of hours staffing for those operations, etc.

The cost of the system, staff support of the system, should be dwarfed by the cost to the business of service loss

Or are you paying to support a cheap system with excessive, expensive staff time and costly outages?

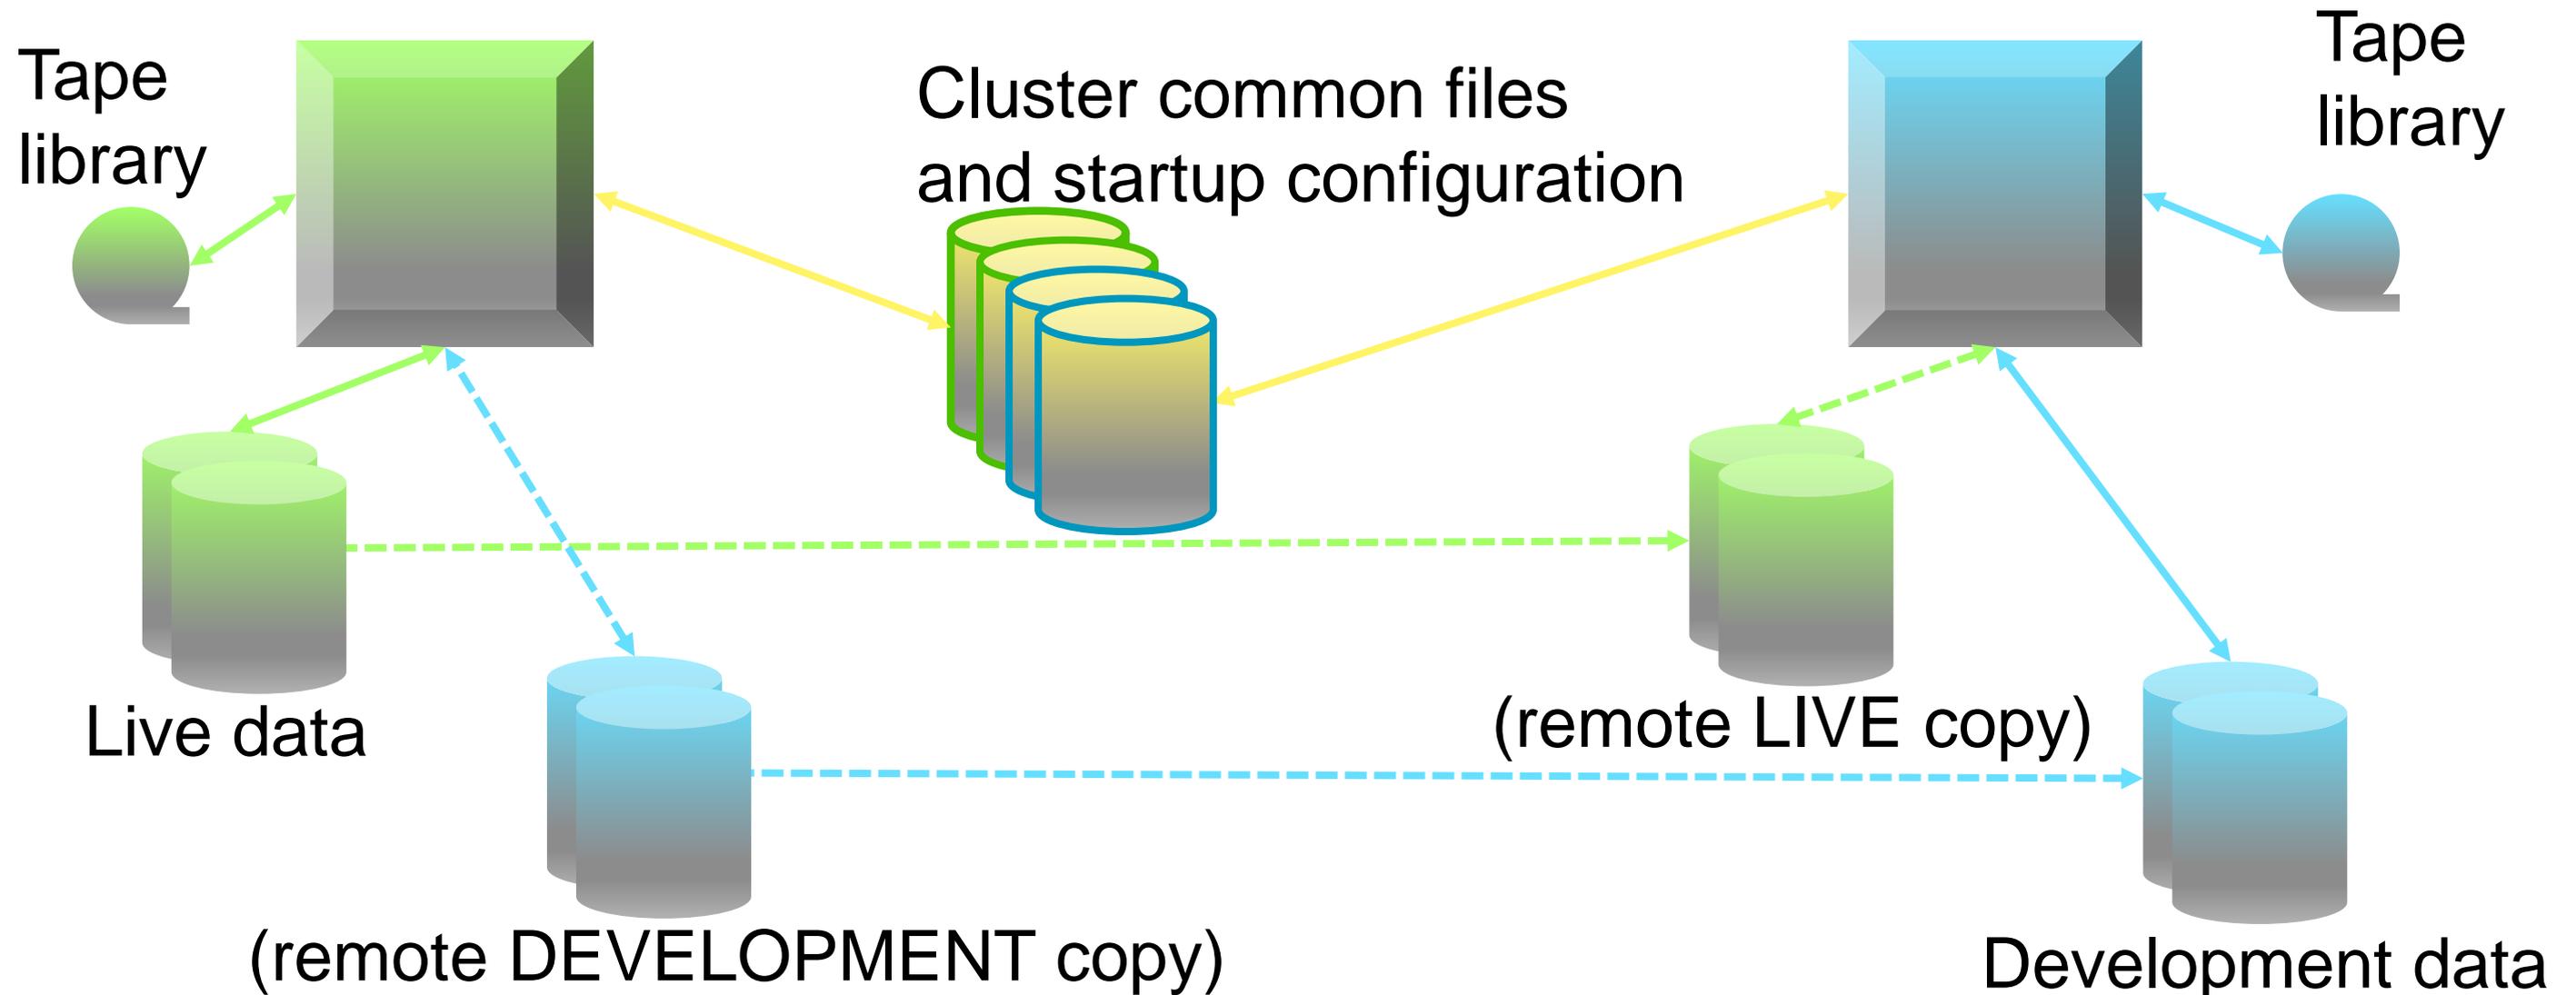
Delivered environment

- Split site IP cluster on OpenVMS 8.4-2L1 with volume shadowing
- Four member shadow sets (two local, two remote)
- At production site-
 - Local mounted production data (copy on development remote)
 - Remotely mounted development data
- At development site-
 - Local mounted development data (copy on production remote)
 - Remotely mounted production data
 - Legacy tape unit for retained backups
- Cluster common disk, individual system disks remotely mirrored

Simple diagram (networks omitted)

Production SITE A

Development SITE B



Interconnections

- Inter-site dual DWDM (Dense Wavelength Division Multiplex)
- Local DNS and gateway (manual change for LIVE service)
 - IP addresses remain FIXED (subnet) per site
- End users “ask” the DNS for the live service
- iLO on internal support network for each server
- No “traditional” SCS
 - UDP / IP instead of proprietary ethernet datagrams

Goalpost moving, or “can you just...”

- IP only (including clustering)
- SFTP and SSH ‘telnet’ sessions
 - FTP and unencrypted telnet **BLOCKED**
- No dial-out modems
 - Move to email SMS responder
 - More informative messages could be sent
- Printer servers (print queues) changed

Methodology

- Cluster built at hardware vendor site, tested
- VLAN between sites for cluster data
- Cluster common disk holding necessary configuration for both PRODUCTION and DEVELOPMENT service
- Booting as PRODUCTION or DEVELOPMENT would read the correct configuration
 - Includes location specific data, such as IP subnets and hardware details which remain fixed to site
 - Network team need to “swing” production DNS to correct site

Technical configuration

- Production system with 2 votes
- Development system with 1 vote
 - Loss of development vote has no effect, QUORUM=2
- Cluster common disk mounted by both systems
- Location dependent start up files (serial number based)
 - Different voting, differing IP configuration
 - But IP routing is per site not per server
 - Redirection for service handled by routers (DNS)
 - 2 local shadowed disks, 2 remote shadowed disks

Testing phase

- Added client network to company firewall rules
- Data SFTP copied onto systems
- Access tests passed
- Batch tests proving far too slow (job comparisons passed)
- T4 used to diagnose, assistance from VSI
 - BACKUP to disk phase in particular very slow (remote writing)
 - Employed BACKUP/COMPRESS to disk
 - Indexed data files tuned
 - Dual network interconnect not correctly working
 - VSI assisted with performance analysis
- Full DR tests, fail over and fail back (ongoing)
- Data copy dry runs

Improvements delivered

- Apart from the obvious (cue sketch from Monty Python's LOB)
- Supported and modernised
- Daily tape backups now completed including Friday
- Full backups now include a set of daily backups, plus images
 - Recovery of a daily is the same regardless of daily or weekly tape
 - Tape backups NOT compressed (drive already in compression)
- Much faster processing
 - Therefore the “out of hours” duty is reduced from “until the early hours” to “before bedtime”
- Alert messages much more detailed

Situations handled

- Tape drive issue
 - Backups made to remote site using TMSCP (remote tape)
- Disk drive issue (failure)
 - Still 3 copies
 - Spares available as part of the design
 - Replaced and merged back during production time



Questions



(This slide intentionally left blank)