

# Re: VMS SAN Primer

---

*Source:* <http://unix.derkeiler.com/Newsgroups/comp.os.vms/2008-07/msg00139.html>

---

- *From:* [etmsreec@xxxxxxxxxxxx](mailto:etmsreec@xxxxxxxxxxxx)
  - *Date:* Mon, 7 Jul 2008 05:07:51 -0700 (PDT)
- 

On 6 Jul, 15:49, IanMiller <g...@xxxxxxx> wrote:

On Jul 6, 3:31 am, Michael Austin <maus...@xxxxxxxxxxxxxxxxxxxxxx> wrote:

David J Dachtera wrote:

JF Mezei wrote:

Paul Lentz wrote:

I sorta knew there couldn't  
be much difference...

But wait a minute, don't SANs use very  
different terminology. They talk  
about switches, fabric etc .

True. However, the terminolgy has become very confused  
(confusing).

When folks say "SAN", they really mean "storage array".

When folks say "fibre channel", they really mean "storage  
area network"  
(SAN – as in the interconnecting infrastructure).

Re: VMS SAN Primer

...and "a separate 'fabric'" equates roughly to a VSAN (Virtual Storage Area Network), corrolary to a VLAN. VSANs taking "zoning" to another level, as it were. On CI or over Ethernet, the VMS equivalent would be the cluster id. number.

And don't SANs have many many capabilities such as RAID, abilities to comvine physical disks into a single drive, or partition a single drive into multiple drives ?

Yes. Think: "HSG" or SWXCR.

Do SANs provide any concept of shared locking ?

Does a CI provide such a concept? ...shared SCSI...?

Can a node request that a block on a drive be locked for writes by other nodes ?

Within the confines of an operating "domain" such as a VMS cluster, certainly. However, it requires a distributed lock manager.

Or is it pretty much a total free for all with SANs just blindly executing requests on any drive from any node ?

Re: VMS SAN Primer

In so far as "drive" and "node" are virtual concepts, yes.  
However,  
there is no "magic" which enables sharing. Read on...

(I would assume that SANs would have  
ability to provide "views" which  
means that a particular node would have a  
defined list of disks it can  
access ?

Yes and no. "LUNs" (remember: FC is just a way to carry  
the SCSI  
protocol over a light "beam") are "mapped" to specific fibre  
adapters  
("FA" for short, in the parlance) on the storage array, and  
"masked" for  
access by specific HBAs (by WWID).

Or can it go and peek at disk drives that have  
been assigned to  
other nodes ?

Zoning, mapping and masking restrict "visibility" between  
specific HBAs  
and LUNs.

Seems to me that there would be a large  
number of management issues to  
deal with that would not be needed in case of  
a VMS cluster. A VMS  
cluster offers a single security concept,  
shared locking etc. When you  
have different separate nodes accessing drives  
in a SAN, those are no  
longer applicable.

Well, you're confusing SANs with MSCP-served storage.

## Re: VMS SAN Primer

The best way to think of a storage array is as if a tremendously talented SWXCR were housed in a rack/frame with fairly large number of physical drives. The physical drives are grouped together by the array manager (a person, that is) into virtual devices. Think: RAIDsets, mirrored RAIDsets (5+1 for example) and mirrored stripe sets. Quite literally, a superset of what's available on an HSJ, HSZ or HSG. Those virtual devices are then presented to specific hosts via zoning, mapping and masking.

...however, it is just storage. A LUN. It's still up to the host operating environment to manage that storage. Such management is NOT the array's job in a FCSF/SAN anymore than it would be in an HSJ on a CI-based storage array. The array simply presents storage. Each LUN appears to the host as if it were a separate "SCSI" device. A "LUN" may occupy a portion of each disk in a disk group (in EVA parlance), for example. VMS, Windows, UX, AIX, etc. only "sees" a SCSI device over FC (\$1\$DGAxxxxx:), while the actual storage presented may consist of a RAIDset or a stripeset, with or without mirroring (on the array, not HBVS).

There's no "magic" in a FCSF SAN which can allow incompatible operating environments to either co-exist or share storage devices. The limitations of each operating environment transcend the storage domain, regardless.

Clear as mud, eh?

Re: VMS SAN Primer

Thought so...

D.J.D.

I took a job a few years ago doing Sysadmin on OpenVMS on SAN.  
Looking  
at the SAN – it is essentially a smart Star–coupler. It directs traffic  
to only those arrays and "devices (aka LUN)" you have specified. Once  
the pointers are set – it is very very easy...

The SAN Switches make up a fabric – which is nothing more than a fiber  
network. Your HBA's can attach to the same fabric – or redundant  
fabrics. Blue/Red for example.

This is an over–simplification, but hopefully you get the idea...

Read the presentation on Data networking and Storage Networking on the  
XDelta site

<http://www.hpug.xdelta.co.uk/>– Hide quoted text –

– Show quoted text –

That doesn't make any note of any of the MSA family of storage  
arrays. Most of the VMS installations that I've seen/been involved  
with use them rather than EVAs, but then maybe I deal with smaller  
clients than XDelta?

.