

## Re: Virtualized VMS in clusters (general questions)

---

*Source:* <http://unix.derkeiler.com/Newsgroups/comp.os.vms/2008-06/msg00815.html>

---

- *From:* Jim Mehlhop <[mjim1025@xxxxxxxxxx](mailto:mjim1025@xxxxxxxxxx)>
  - *Date:* Sun, 22 Jun 2008 09:44:06 -0600
- 

Alan Winston – SSRL Central Computing wrote:

VMSers --

I'm trying to wrap my head around how virtualized VMS systems participate in certain aspects of clustering and volume shadowing. There may be something I'm just not getting. So this is kind of general.

At the HP Tech Forum I got to play with booting a virtual Itanium on a real Itanium blade (which ran HPVM which runs on HP/UX; my first HP/UX login ever). I've encountered, also, the SRI Alpha emulators, and seen the SimH VAX emulator running. (I was trying to ask the kind of question I'm asking here at the Wilm's very interesting session on the architecture of the Alpha emulator, and I made them sound like plain VMS questions, but they're really questions about the interaction of the emulated VMS node with the cluster, and I didn't formulate them very well in person. I totally get that what the SRI Alpha emulator provides is an Alpha inside your Windows box, and that once it's up, VMS is VMS – or Tru-64 is Tru-64, or Linux is Linux, or whatever.)

I get that if you're standing outside a box running a VAX or Alpha emulator under some other host operating system, you might as well be standing outside a VAX or Alpha. The virtual machine interacts over the Ethernet just like a real machine, and if you can keep the interaction to the Ethernet, you've just got a fast cheap VAX or Alpha, or more than one. (I've seen two SimH-emulated VAXes on the same laptop clustered together over Ethernet, for example.)

The host operating system forwards Ethernet traffic to a virtualized NIC on the virtual VAX or Alpha, which can then participate in an NI cluster, no problem. (Or maybe the host is able to have multiple NICs and dedicate one to the virtual machine.)

But I'm curious about access to shared resources in other ways. I think the relatively trivial answer is to run your virtual nodes as satellite nodes over Ethernet, and then everything just works. (Assuming that the host OSes appropriately forward MOP requests and the responses thereto. Don't know if the host OSes care about protocols or if they'll just forward anything intended for that MAC address.)

How does access to SAN disks work? (If my real host is Windows or HP/UX or

Re: Virtualized VMS in clusters (general questions)

Linux, and the real host has a Fibre Channel connection, then I have a box that isn't participating in my Distributed Lock Manager that of necessity has write access to my cluster disks, which seems like a bad idea. [I mean, the real host has a wwid and the EVA has to present the disk to the real host, right?] Is there some way in which I can dedicate a real host Fibre Channel connection to my virtual machines? Do I have to have one real VMS box with a Fibre Channel connection on the cluster presenting my SAN disks to the virtualized hosts over MSCP?

This one I would worry about

What about single-system-disk clusters? Virtualized VMSES don't actually know what underlying device they've booted from, and they might not actually have booted from the same device. (You present the disk image, which might be a container file, a DVD, or whatever, and it looks to the emulator like a generic disk – in the hpvm case, a generic SCSI disk with the unit number you give it.) Can multiple virtual VMSES boot from and log to (do SYSUAF updates, put stuff in a cluster wide audit log, etc) a cluster-common disk?

YES I do this with Personal Alpha all the time

Who

arbitrates access to it? Can you do shadowed system disk on it?

I would say yes but I have not tried it.

How?

Is an all-virtual cluster that uses the single-system-disk approach possible?

yes

On the Itanium-emulated-on-Itanium approach (which isn't supported until VMS 8.4) can you successfully run a multi-node cluster all on one physical box with a single system disk?

Again I would say yes but without testing, it is my opinion only

[There's some appeal to this one because virtual

machines inside the box can communicate via a virtual switch, with traffic never leaving the box – faster, no packet sniffing, etc. But I'm not clear

Re: Virtualized VMS in clusters (general questions)

Re: Virtualized VMS in clusters (general questions)

whether the hp/ux virtual switch supports, eg, DECnet, MOP, SCS packets (although I note that 8.4 also contains clustering over IP, so if the virtual switch doesn't support SCS-qua-SCS but does support IP (which it \*must\*), this could still work.)

Can you do host-based volume-shadowing on SAN disks that are actually being presented by non-VMS hosts (if that's how that works at all)?

Again I am not sure this is possible

How about on

container files presented by the host?

Yes as an example of this use the lddriver that is included in VMS 7.3-2 and beyond for testing. It is a container file just like the virtual disks from SIMH or personal alpha

(And in that case, does the

enterprise just rely on system managers being very careful to use the same unit numbers for the same files everywhere, because VMS just doesn't have the information to let you know you screwed up?)

The same rule for volume shadowing apply if you are using emulated virtual disks as normal disk. must have enabled and licensed the shadowing software and the "device" must be presented with an allocation #

Any insight appreciated.

Thanks,

-- Alan