

Re: Blast from the 1988s (DEC proposal)

Source: <http://unix.derkeiler.com/Newsgroups/comp.os.vms/2007-01/msg00832.html>

- *From:* etmsreec@xxxxxxxxxxxx
 - *Date:* 16 Jan 2007 05:21:24 -0800
-

I think you're right Mark – the message I've heard from a number of people is that it's a case of needing to take lots of peoples' workstations over for a weekend in order to gather enough systems together to form a cluster that big – most companies only have that number of PCs available for this kind of job. Few have that many VMS systems available.

The limit probably got changed with the introduction of LAVc.

Steve

Mark Daniel wrote:

Mark Daniel wrote:

Larry Kilgallen wrote:

In article <87ac0kdrjo.fsf@xxxxxxxxxxxxxxxxxxxxxxxxxxxx>, prep@xxxxxxxxxxxxxxxxxxxxxxxxxxxx writes:

JF Mezei
<jfmezei.spamnot@xxxxxxxxxxxxxxxxxxxx>
writes:

At the time, the graphs showed from the all mighty microvax II to the VAX 8978. The document mentions up to 15 8800 nodes in a cluster. (when did it go from 15 to 96 nodes in a cluster ?)

Re: Blast from the 1988s (DEC proposal)

AIR, the 15* was the limit on the number of nodes on a single system volume, not the cluster size limit. I *THINK* that back then the size of a cluster was limited by the size of the statically allocated System Director Vector for the DLM.

Formerly there was a distinction between the total number of nodes supported in a cluster and the number of large nodes supported in a cluster. This was based on what had been tested, not theoretical limits.

From memory (and it's been a while since I've actually seen a CI – though I think we've still got the cables between our two, redundant machine rooms); it was 32 nodes per CI (the star-coupler), half of which could be VMS systems and the other half storage controllers (e.g. HSCs). The number of CI's (redundant paths or independent storage) on any one system depended on the size of the iron. We (HFRD->WASD->SSD->ISRD) had over the span in excess of a decade a time-variable collection of systems based on a mixed-interconnect cluster, including a handful of CI-based systems (including at one stage a VAX9000), along with NI-based VAXservers, VAXstations, DECstations, AlphaServers, AlphaStations, etc. Our main cluster exceeded 70 systems at one stage.

What I would have added had I not inadvertently hit [send]; IIRC the 96 nodes was a supported maximum because that was the largest cluster Engineering had managed to configure on the test-bench. I'm (idly) curious about design limitations (ignoring the practical limitations of memory, bandwidth, etc.) on cluster size. Things like the 65k DECnet node limit. (Hope the spelling's less objectionable this time, Steve.)