

Re: Page datasets on SAN

Source: <http://unix.derkeiler.com/Mailing-Lists/AIX-L/2005-06/0166.html>

From: Damir Delija (*ddelija_at_PUBLIC.SRCE.HR*)

Date: 06/26/05

Date: Sun, 26 Jun 2005 06:02:47 +0200
To: aix-l@Princeton.EDU

On Sat, Jun 25, 2005 at 04:05:26PM -0700, Sandor W. Sklar wrote:

- > *Well, according to the EMC/IBM support matrix, booting AIX systems*
- > *from Symmetrix or Clariion disk is supported, but I have no*
- > *experience doing that. It does seem a bit of a catch-22, being that*
- > *you HAVE to have a paging space in the rootvg, and yet EMC recommends*
- > *against putting AIX paging space on SAN devices. Since I have no*
- > *plan on booting the AIX systems from the SAN, I've never forced*
- > *through EMC support a reconciliation of the apparent contradiction.*

Actually I've done very intense testing on LPARised p690 under AIX 5.2 ML0 with swap on symetrix and it worked perfectly during test in all possible combinations.

But later during production due to misconfigured SAN we get most of the problems. One cause of production problems was lack of powerpath during boot which causes panics this was done by our SAN people against IBM advice. When all inidents were analysed it comes down to faulty desing and lack of knowledge symetrics and AIX wa dooing well people were doing mistakes.

We also look about SAN boot but in there were some reasons against in case of SAN failure it will be hard to do anything I was really scared when I saw literally hundreds of devices as possilbe boot disks in SMS menus on regatta

My feeling is that SAN makes things complpated for a level of magniuted and this cause trouble.

I hope this helps

Damir

- >
- > *We used to boot all of our Solaris systems from the Symmetrix, but as*
- > *we moved systems to Clariion, we ran into lots of problems, and it*
- > *just seemed as if it was more of a headache then anything else.*

AIX-L: Re: Page datasets on SAN

> *We're now booting our Sun systems off of internal disk, mirrored via
> DiskSuite/SVM.*
>
> *I'd be interested in hearing the reasons why your co-worker doesn't
> like the IBM rootvg mirroring system; in my experience, I've yet to
> see any other OS do it better, either with built in tools or add on
> software.*
>
> *Good luck!*
> *-s-*
>
> *On Jun 24, 2005, at 10:13 AM, Vincent D'Antonio wrote:*
>
> *>based on this one would say you should never put your rootvg out
> >there? Someone here not an AIX person doesn't like the way IBM does
> >it mirroring and want to move the rootvg to the EMC disk.*
> >
> *>any thought's.*
> *>Vince*
> >
> *>On 6/24/05, Sandor W. Sklar <ssklar@stanford.edu> wrote:*
> >
> *>>We've got sort of the same problem on this issue. I've never seen an
> >>IBM document talking about this issue, but we use EMC PowerPath (with
> >>our Clariion-based SAN), and in the PowerPath 4.4 for AIX Release
> >>Notes document, it states this (on page 24):*
> >>
> *>> In rare cases, if the AIX swap device is stored on a CLARiiON
> >> LUN, your system may hang if all paths to the active SP are lost
> >> and the AIX code, required for PowerPath to complete a trespass
> >> operation, is swapped out of memory. If this problem occurs, you
> >> must manually trespass the LUN to the available SP and reboot
> >> the host.*
> >>
> *>>Unfortunately, aside from two internal drives (mirrored rootvg), the
> >>only storage on the system is Clariion, so we've had to with paging
> >>spaces on the SAN.*
> >>
> *>>We have had occasions where some of our AIX systems have hung
> >>(requiring hard reboot) due to a trespass issue on the SAN, but I've
> >>never been able to track down the definitive cause.*
> >>
> *>> -s-*
> >>
> >>
> *>>On Jun 24, 2005, at 7:17 AM, Holger.VanKoll@SWISSCOM.COM wrote:*
> >>
> >>
> *>>>you are talking about paging-spaces, dont you?*
> >>>
> *>>>I've never seen a doc from ibm that states ps on san is not*

AIX-L: Re: Page datasets on SAN

> >>>recommended and would be surprised if you could provide a link to
> >>>it.
> >>>
> >>>if "ibm says" means "somebody thats working at ibm told me..." then
> >>>you can safely ignore it. you can find others at ibm that tell the
> >>>same, and others at ibm that tell something different.
> >>>
> >>>there are ps-recommendations; in short
> >>>dont place a ps on a disk where a ps already exists
> >>>dont place a ps on a disk where a mirror of a ps already exists
> >>>
> >>>and, by far not so important
> >>>keep all ps at the same size
> >>>
> >>>the only argument against ps on san that I see is cost
> >>>if your san is less reliable than internal disks, also reliability
> >>>is of concern – if a ps is not accessible anymore, the system will
> >>>crash
> >>>
> >>>
> >>>
> >>>From: IBM AIX Discussion List [mailto:aix-l@Princeton.EDU] On
> >>>Behalf Of Morris, Carey B
> >>>Sent: Friday, June 24, 2005 3:16 PM
> >>>To: aix-l@Princeton.EDU
> >>>Subject: Page datasets on SAN
> >>>
> >>>Is anyone aware of any good white papers or redbooks regarding
> >>>placing AIX page datasets on SAN storage? IBM says it is supported
> >>>but not recommended. What is their reasoning? Performance?
> >>>Reliability?
> >>>
> >>>Thanks,
> >>>Carey Morris
> >>>
> >>>
> >>>
> >>>