

Re: Exchange 2007 – Multiple SG per LUN?

Source:

<http://www.tech-archive.net/Archive/Exchange/microsoft.public.exchange.admin/2007-12/msg00160.html>

- *From:* "Mark Arnold [MVP]" <mark@xxxxxxxx>
 - *Date:* Mon, 03 Dec 2007 17:04:43 -0500
-

Sorry I meant RAID 1+0!!!! and yes I know what a LUN is :)

It's DAS/HP MSA array. I was just wondering why best practice appears to of changed. I would of normally seperated each SG onto its own RAID 1+0 LUN for max performance and also have the transaction logs (RAID1) on their own LUNs seeing as we will have a small number of storage groups. I can understand why combining them would be beneficial on a much larger server housing a large number of storage groups. We were planning on 4 storage groups each with 1 store (following BP and because its a CCR server). With a 4 drive minimum for RAID 1+0 that would be 16 drives (for the stores) and this configuration would provide enough IOPS and capacity to support the user base. For the transaction logs we were planning on 4 RAID1 LUNS making a total of 24 disks.

Now why would combining some of these stores on one LUN be beneficial especially seeing as the disks have already been bought!?

From your reply, I think you are agreeing with my view point?

Thanks

AJ

So the MSA is just a little bit o' DAS. You haven't got enough disks to implement a "plus" anything in RAID so I'd stick with RAID1 and call it right.

Is this MSA tied to the network or have you got two of them? With that system (dreadful throughput) I wouldn't pay a terrific amount of attention to what MS say in the calculator. It's not a good system in terms of I/O capabilities. RAID1 for a couple of SG logs and then RAID1 for each of the stores you're planning. So that's eight disks

Re: Exchange 2007 – Multiple SG per LUN?

for stores and then another 4/6/8 for logs and another pair for queues and queue DB; 18 disks so far for one node.

I have a low opinion of the benefits of RAID1+0 or 0+1 really. They don't achieve a great deal and, given that people who use that configuration will tend to stick the disks that came (i.e. same batch) together you're every bit as likely to suffer one disk failure as two (I exaggerate slightly for demonstrative purposes but you get my point).

.