

Re: Disk Array Usage

Source: <http://unix.derkeiler.com/Newsgroups/comp.sys.sun.hardware/2004-02/0106.html>

From: Paul S. Brown (usenet060301_at_geekstuff.co.uk)

Date: 02/06/04

Date: Thu, 05 Feb 2004 23:33:46 +0000

Robert Mazur wrote:

> *I have a question about the usage of disk arrays. I have not
> had the chance/need for building a "larger" system. I admin
> about 6 Dell Poweredge servers running Linux. One is a database
> server, and the rest are web, file and app servers.
>
> My Sun/Solaris usage is limited at the moment to using my U60 and U5
> as desktops for daily usage.
>
> For my own education, how does using a disk array differ from just
> using a server in its place? I see an array as a massive storage
> facility, whether that be a file server or running a database. Rather
> than a disk array, why not have an actual server with a ton of space?
> Either way you can have multiple web/app servers accessing the data?
>
> Could this be because such a server with high storage capacity would
> likely have high processing power, and therefore you are paying for
> power rather than storage?
>
> Are all disk arrays connected to their servers via Fiber Channel
> cables,
> making for better through-put?
>
> Can someone explain to me where disk arrays fit into the larger
> system configurations?
>
> Thanks for your help!
>
> p.s. I am aiming at our next server purchase to be a simple Sun
> box like a V120.....just because! :-)*

I think you need to redefine your terminology here.

A disk array is simply a lot of disks accessible by a system. This can be as simple as the built in disks on your poweredges. If you have a disk array which can be attached to only one machine then this is normally known as "DAS" or Direct Attached Storage.

If you have a number of disks which can be attached to by several machines at once then this is normally a SAN or Storage Area Network. All this means is that several machines can access the disks at once. Normally this access is by means of FC-AL which is Fibre Channel Arbitrated Loop. This is a means of attaching lots of disk controllers in a chain and allowing any unit on the chain to talk to any other. Units can either be storage devices (disk drives) or hosts (servers). In higher end SANs the access is via switched Fibre Channel which works much like switched ethernet in that any device can ask to speak to any other device and they will have a dedicated path between them for the duration of the communication. This switching is normally achieved by Director Switches.

If you have a number of disks which can be accessed over a TCP network then this is normally known as NAS or Network Attached Storage. This normally has a head unit attached to some kind of SAN with the head unit providing network, NFS, CIFS and the like. In the case of NetApp the head unit is effectively a high spec PC running an OS (DataOnTap) which is optimised to hell and beyond for network operation. In the case of BlueArc most of the services are performed in hardware.

You seem to be asking about NAS storage and the answer is that there is very little difference between a NAS filer and a server with a ton of disk attached to it. The only thing the filer tends to have going for it is that its software is better optimised for the duty than a general purpose OS. This is something of a generalisation that only actually applies at the high end, but it's a useful idea.

SAN storage tends to be attached via FC running at either 1Gb/s or 2Gb/s. The main thing that FC has going for it is that it has minimal overhead as all it transfers is basic SCSI III commands and a few control commands. Nowadays there is a concept coming up called iSCSI which is blurring the lines. Where FC is SCSI running over dedicated fibre iSCSI is SCSI encapsulated in TCP which allows you to use your existing IP network for your storage infrastructure.

Where SANs start becoming useful is when you have lots of machines with varying disk space requirements. You could just attach storage directly to each machine, but this is very inefficient as you can end up with machines with not enough disk space and machines with far more than they can use and no way to exchange the space. With a SAN you can dynamically allocate the space so that each machine has the storage it needs and no more nor less. You also get the ability for several machines to attach to the same logical disk which allows for such things as failover of filesystems between machines and application clustering such as Oracle 9i RAC provides.

SANs are basically complicated networks of disks and controllers with some intelligence built in. Most SAN systems provide support in the hardware for RAID – with EMC kit this tends to be RAID 1 (mirrored disks) which is as resilient as you can get. HDS kit supports RAID 0,1 and 5 which gives you an option (never run critical applications on RAID0 – it's fast, but if you lose one disk you have lost all of your data).

P.