

ZFS BASED STORAGE SERVER

Using Solaris 11 EXP, applies to Debian/BSD, ION, OpenSolaris

OVERVIEW

- Storage Server/Home Storage Server Concepts
 - What is a Home Server
 - What can I do with a Home Server
 - What kind of connectivity will my Storage Server have
 - ZFS Primer
 - What is ZFS and why do I care
 - Containers
 - Concept of Containers vs Virtualization
 - Hardware
 - SSD's
 - Storage Controllers
 - Processing Power
 - Software
 - What you will need
 - Basic configuration guidelines
 - Gotcha's
-

EDITORIAL

- There are a number of how-to's out there using ZFS as well as OpenFiler, Freenas, Nexenta and others as a basis for a home server. I will only touch on those briefly during the course of this presentation, I've installed them all tried them all in both home and semi -production environments and will share what I've discovered with you in future articles for our LUG site. I've made a significant number of personal discoveries during this process, not all of them good. I will have to say we will not be able to cover all the deep dive details during one session.
- This set of slides and articles will address both the enthusiast, and the professional as I am (some would say arguably) both. You will see references to technologies the home user will most likely not have access to and some references to processes the professional will not really care about, I would encourage you all to find our LUG site, sign up for the mailing list and comment anyway even if it has no bearing on what you want this device to do for you.
- I am by no means a Solaris expert (as you will discover) so please be kind when you post about how retarded some of the methods I use are. I use primarily Linux and most of you at our LUG will probably be wondering why Solaris, and I'll answer that question in the course of the presentation. I started this project hell bent on using
- FreeNAS btw, not Solaris.

STORAGE SERVER/HOME STORAGE SERVER CONCEPTS

- The idea of a home server is a central place to store files, movies music and just about anything that you hold dear to your heart.
- A central point of administration, backup and permissioning
- One can have their server be a media server offering streaming video and audio to multiple points throughout their house and with the use of some slick networking anywhere on the internet.
- The Storage server can also serve as a central point for server virtualization as well as a point to protect all of the important files that live on the laptops and desktops in your home or office.
- The basic type of connectivity in a storage server is File based using either NFS, or CIFS/SMB. Almost all the current generation of storage server software also allows for Block storage access ie iSCSI or FCoE even Fiber Channel. In the Home the most common connectivity choice is going to be Ethernet taking advantage of File based sharing technologies found on most windows and mac based PC's.

WHAT IS ZFS AND WHY DO I CARE

- Zettabyte File System (ZFS) is a mature file system developed as a copy on write (cow) file system, meaning the data on disk is not changed it is copied to an overflow area of the disk before the change is committed maintaining data integrity even when things don't go exactly as planned. COW is not unique to ZFS, WAFL (NetAPP), BTRFS and ext3COW are all examples of other copy on write file systems. Windows has a cow style feature called volume shadow copy but NTFS is not a cow file system.
- ZFS does away with the need for hardware RAID controllers. It manages the hardware itself using what are called Pools or zPools which are made up of physical disks.
- Huge limits, I mean ... big huge limits. The limits for the number of files in a folder and the number of folders etc are astronomical, ZFS is a 128bit file system (So is BTRFS from what I understand), meaning you are going to run out of disk blocks before it runs out of inodes. Does it matter? Absolutely, if you have all your music, word documents etc for your whole family and your work stuff on one volume in one folder tree starting from / (root) you WILL run out of inodes using EXT2 and some other older file systems.
- The data integrity processes built into ZFS exceed what most people will ever find themselves needing. If you are like me and make mistakes (well I thought I made a mistake once, turns out I was wrong) or if you have power outages or a mix of bad environmental and bad equipment ZFS does a great job of preserving your data.

THE IDEA OF CONTAINERS

- Solaris has the concept of “Containers” vs the idea of a Virtual machine. In theory they are almost identical with the container being a virtualized instance running within the same operating system and a virtual machine being a completely separate instance.
- This is where we run Debian! Some applications are not quite yet available on Solaris and others are just much easier to use in Linux.
- DLNA, UPNP, Itunes and Apple streaming is facilitated by running a Linux instance within a container.

HARDWARE

- While none of these things are written in stone, I would suggest starting with at least this and going up from there
- Computer with at least 8GB RAM and 2GHZ processor (multi-core preferable)
- At least 3 hard drives of similar capacity in JBOD (Just a Bunch Of Drives) mode NO RAID!!
- At least 1 network card running at least 100mb/s preferably 2 cards if you have 100mb/s.

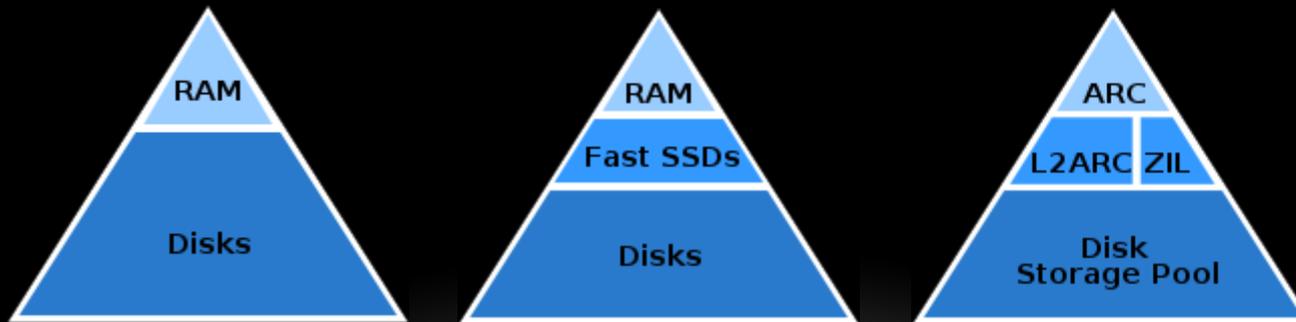
This is the basic hardware needed for a home server. Obviously you can get a little fancier with hot swap drive bays, multiple controller cards for up to 16 hard drives 6, 8 or even 12 core processors but that will be defined by your budget and what you intend to do with your system.

If you plan to do many things like Torrents, file serving, iSCSI etc then you should look at least a quad core 3ghz cpu and 8GB RAM at the absolute minimum.

HARDWARE CTD

SSD's can be used to significantly speed up the performance of your ZFS based system by using them as Log disks, ZFS Intent Log (ZIL) cache, and L2ARC (Layer 2 Adaptive Read Cache).

- A 16GB SSD used for a ZIL cache is plenty for all but the most demanding of home work loads, more is obviously better when it comes to SSD's but ... 16GB is plenty.
- 72GB+ for the L2ARC, while I am a little hesitant to say this is end all be all of recommendations but based on my research and hands on with a retail storage array 72GB seems to be a decent amount for a read cache.



HARDWARE CTD

The onboard Intel and or AMD chipsets on most boards come with at least 4 SATA ports some of the enthusiast boards will have up to 8 because they use 2 controllers. If you want more than that or you don't have any ports free in your case you can always add a controller card. My experience with this has been to buy cards that don't get in the way but are true hardware ports. I have an Areca 1220 RAID controller in my server configured in JBOD mode but honestly that is overkill, some SiiG or Saitek SATA controllers will make your machines just as nice and just as reliable as any high dollar card.

- Ports are king, don't worry about RAID performance
- Split amongst multiple cards to help mitigate card failures
- Try to use PCI-E if at all possible to alleviate
- potential bottlenecks

HARDWARE CTD

Your processing power will dictate what features of ZFS you can leverage to the full extent. ZFS has:

- Encryption, which can encrypt all at rest data and can be applied per “folder”
- Compression, which is user selectable amongst quite a few different options GZIP and LZJB algorithms
- Deduplication which really stresses RAM along with CPU
- Fiber Channel access, I’ve found my system to really need some horsepower when being accessed via the HBA in the system. When I have multiple things going on I can see the CPU spike vs on a traditional file system I don’t see those loads.

SOFTWARE

I used Solaris 11 Express, but you can use any OS that has ZFS built in like FreeBSD, Opensolaris, ION and Nexenta. FUSE unfortunately will be a very slow performer and doesn't allow you all of the features you get with one of the distro's that have ZFS native to the kernel.

Nappit which is a NAS/SAN web administration kit

Transmission a command line bittorrent client

TFNG TorrentFlux NG is a web based multi user torrent manager

XBMC Xbox media center manages your media and allows for streaming to other machines on your network

Debian/Fedora/Linux distro of your choice for DLNA/UPNP functions.

LINKS

- <http://blogs.sun.com/brendan/entry/test> <= stole the pyramid image from him
- <http://en.wikipedia.org/wiki/ZFS> various limits of the file system
- <http://www.napp-it.org/> nas web admin software
- <http://www.transmissionbt.com/> transmission
- <http://www.torrentflux-ng.org/forum/> torrentflux ng