

Building Hardware to Help You

Or “How to make a big fast Fusion drive
out of 2 RAIDs with SSDs and HDs”
presented by Chris R Tarnowieckyi

Let's review some
technology.

Review of Fusion Drive

- Apple Introduced Fusion drives in 10.8.2
- Fusion drives were built into some models
- SSD combined with a regular HD
- Made possible by software (CoreStorage)
- Advantages in speed and storage size
- Reliability concern, either device could fail
- Supports Recovery HD and FileVault

Fusion Drive, How it Works

- A SSD is tiered with a HD and acts as one device
- SSD boosts read and write performance
- Frequently accessed data is stored on the SSD, less frequently used data on HD
- The system takes care of moving blocks of data as needed to the correct device without the user manually managing the locations of files

Fusion Drive DIY

- You can build your own with “diskutil corestorage”
- Warning: Apple won’t support DIY Fusion drives
- Internet search “DIY Fusion Drive” for plenty of articles
- “Apple Fusion Drive—wait, what? How does this work?” Ars Technica
- “Achieving fusion—with a service training doc, Ars tears open Apple’s Fusion Drive” Ars Technica

Review of Hardware RAID

- Hardware RAIDs look like single block storage devices to the system
- May provide redundancy, speed and large capacity
- Usually supports recovery hd, FileVault, and partitioning
- Apple or third party
- external or internal

What about Apple Software RAID?

- Part of Apple storage system
 - Striped Volume (RAID 0)
 - Mirrored Volume (RAID 1)
 - Concatenated Volume (Spanning)
- Often used as an inexpensive alternative to HW RAIDs
- Doesn't support Recovery HD or partitioning
- I won't use software RAIDs in this project, but...

There is an interesting statement in the diskutil man page

- a PV [Physical Volume that is part of a Fusion drive] is normally real media but it can be a disk image or even an AppleRAID Set.
- What? Does this mean that an Apple software RAID can be a part of a Fusion drive?
- I tried it, I couldn't get it to work. Maybe some other time...




My Vision of a Fusion Drive of RAIDs

- RAID of SSDs for speed and reliability
- RAID of HDs for capacity and reliability
- Fusion drive of RAIDs has a combination of the above features
- Supports all Apple storage features like Recovery HD and FileVault

What should I call a Fusion Drive of RAIDs?

- FDR
- FU-RAID
- Fusion/RAID







Fusion/RAID



Nearly any Mac can have
a Fusion/RAID
Let's dream a little...

The Mac Pro “Cylinder”

- Expansion Busses
 - ThunderBolt 2
 - USB 3
- Built in PCIe Flash Storage
- Gee, I wish I could afford one.



External RAID

- Thunderbolt connection
- Could be a part of a Fusion drive
- Gee, I wish I could afford one.



I can't afford my dream
system, what can I do
right now?

My Mac Pro mid-2012

- Needs “More” for video editing
 - More storage (now)
 - More speed (maybe later)
 - More memory (soon)
- Expansion Busses
 - USB 2
 - FireWire 800
 - PCIe 2.0
- 2 Optical Drive Bays and #2 is empty



My Fusion/RAID

The Parts

- Internal RAID 5 of 4 SSDs attached to an ATTO RAID controller
- Internal RAID 5 of 4 HDs attached to Apple RAID card
- Supports all Apple storage features like Recovery HD and FileVault

A Note about PCIe cards

- Some are only Host Bus Adapters (no RAID)
- Some are RAID controllers
- Some are limited to 3 Gb/s
- Some have SSDs embedded on the cards
- Always check the technical specifications to see if the product meets your requirements.



The HD RAID

- Apple Hardware RAID card for Mac Pro (2009)
- 3 Gb/s throughput per SATA channel
- Four 2 TB drives (limit 2.2 TB)
- Apple Support articles: HT202084 and HT201160
- Note the battery on the card is an item that needs periodic replacement.






HITACHI 2TB 7200 R3.5 HD

These drives were refurbished and are limited to 3 Gb/s.
Such a good price!

The SSD RAID

- ATTO 644R RAID Controller
- 6 Gb/s throughput per SATA channel
- 4 internal / 4 external possible
- four Mercury Electra 240 GB SSD drives
- Will work in other computers and PCIe chassis if I switch from the current Mac Pro





4 Solid-State Drives

240GB Mercury Electra 6G 2.5-inch 7mm SATA 6.0Gb/s



Multi-Bay Internal Enclosure for the 2nd Optical Drive Bay

ICY DOCK MB324SP-B ExpressCage 4x2.5" SAS/SATA
HDD Hot Swap Mobile Rack



Back of the Multi-Bay




Cable connects RAID Controller to 4
SSDs

StarTech Model SAS8087S4100 39.4" (1m) Serial Attached
SCSI SAS Cable - SFF-8087 to 4x Latching SATA



Cable to take power from existing optical
drive cable

SATA 22 Pin Male to 7 Pin SATA Cable with 15 Pin SATA
Female Power Cable



How does it fit?

Especially the 4 SSDs?


The Procedure

- Research, Purchase, wait for deliveries
- Backup Data
- Install Hardware
- Use utilities to create and format the RAID sets separately
- Install OS on RAID sets independently
- Test RAID sets (booting and throughput)
- Create Fusion/RAID drive
- Install OS
- Test Fusion drive (booting and throughput)
- Restore Data

This is the good part.

Install Hardware

- Install the new HDs
- Install the SSDs in the multi-bay enclosure (easy)
- Install the multi-bay enclosure in the 2nd optical bay
- Install the PCI cards (easy)
- Install the cables (challenging)
- Re-install the optical drive bay enclosure



Install the New 2 TB HDs in the Carriers


These carriers are easy to remove; there are 4 screws to affix the 3.5" drives.



Install the SSDs in the Trays

Plastic trays
2 screws secure the SSD to the tray






Empty 2nd Optical Drive Bay

Ready for the Multi Bay ExpressCage



An Enclosure in an Enclosure

secure with screws, but be careful not to drive the screws into the SSDs




Install the PCI card

no cable attached because...



Install the cables


The difficult part






Install the cables

Cables behind the plate





The cables connected to the back of the ExpressCage

To get this picture I had to just let the optical drive bay enclosure “dangle” on the attached cables; not recommended



Re-install the optical drive bays

Done with cables, whew!

This is the better part.

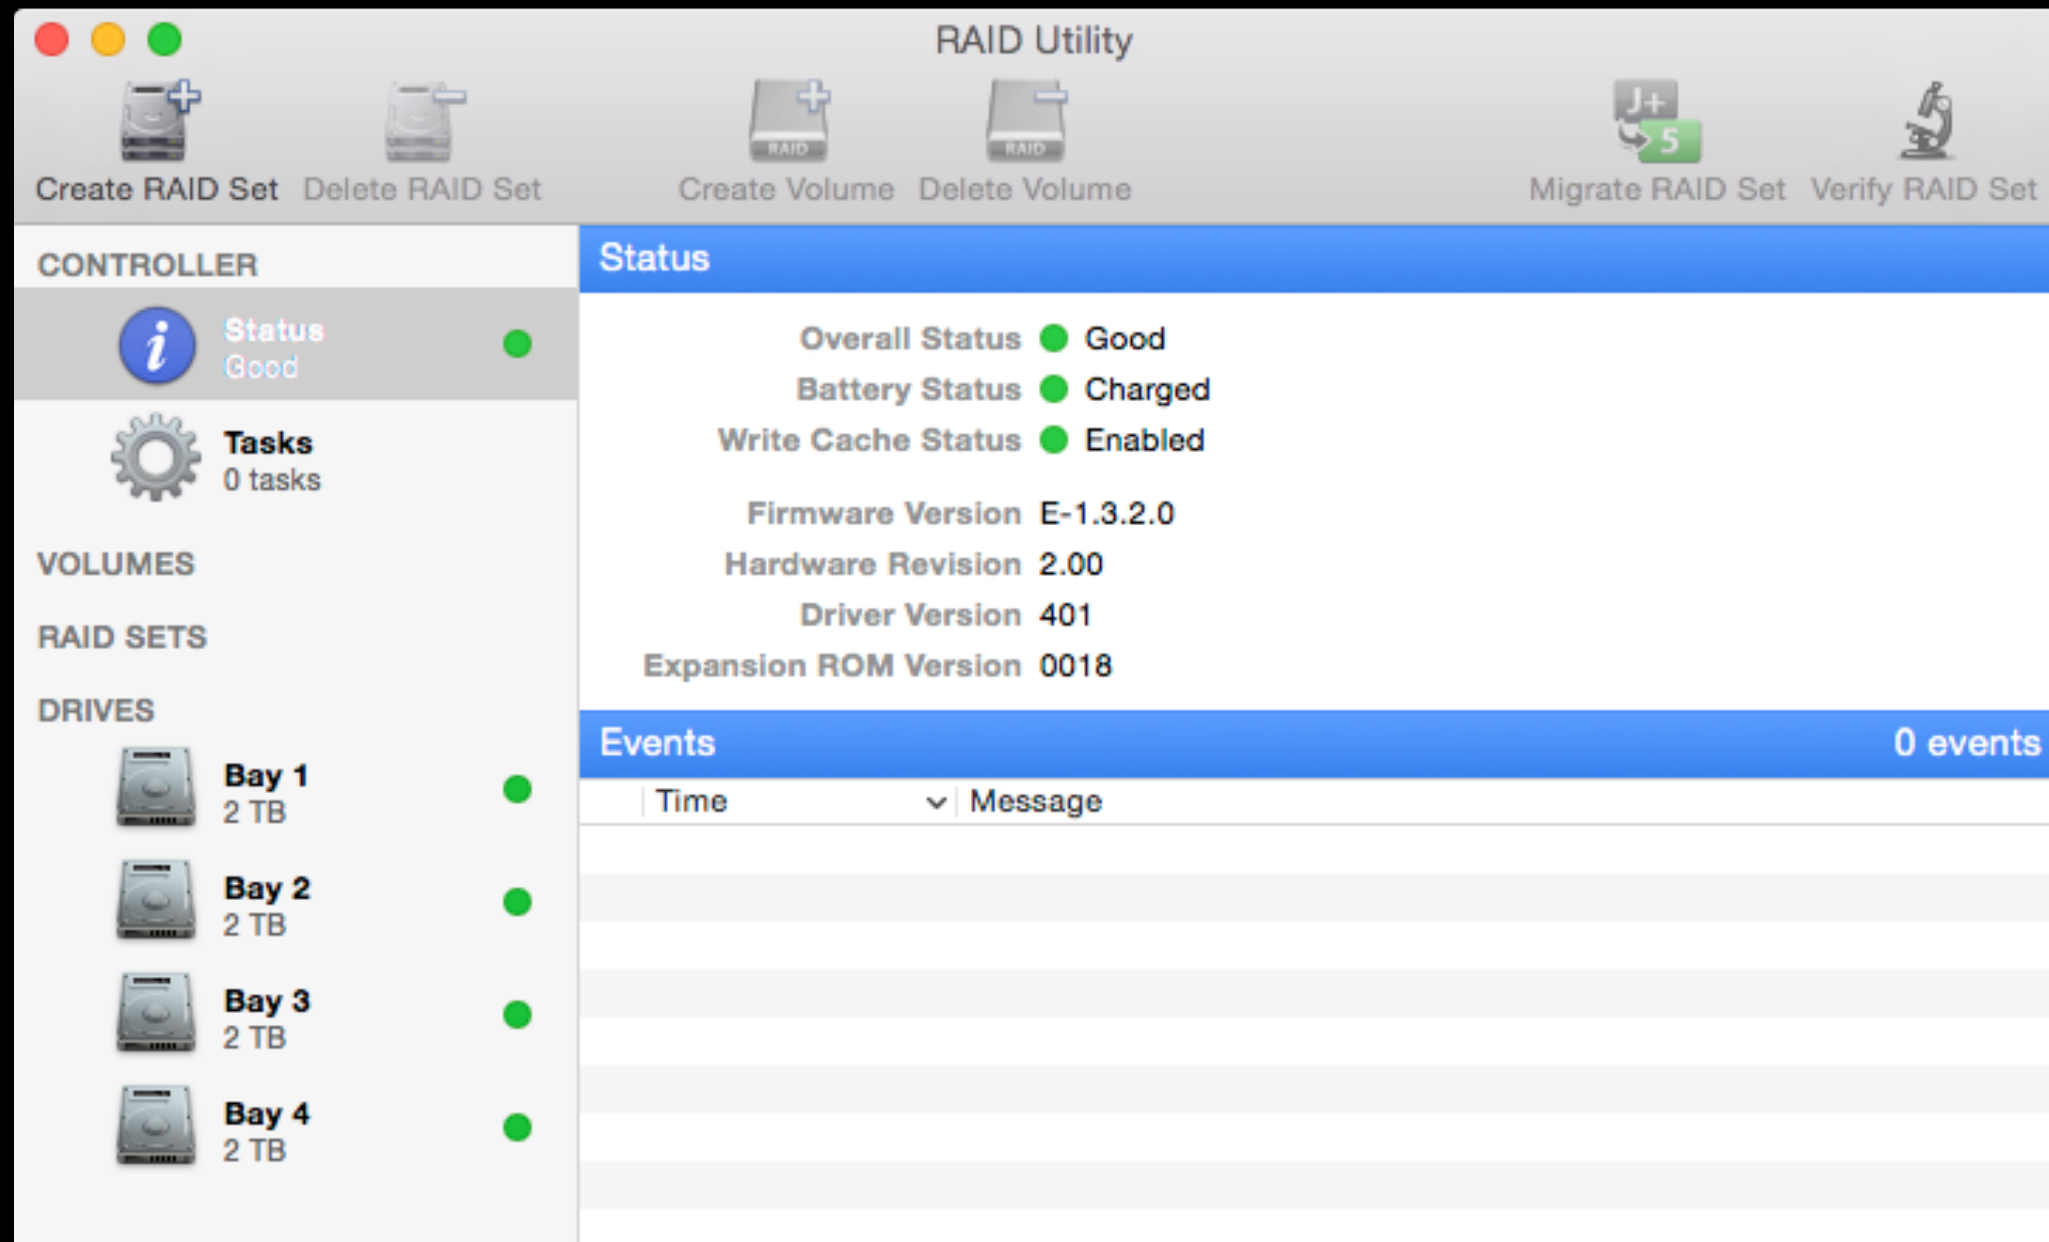
Use Utilities to create RAID sets

- Create the HD RAID
- Create the SSD RAID

Create the HD RAID

- Apple RAID utility is installed with the OS
- Located in /System/Library/CoreServices/Applications
- Creating the RAID 5 set is straightforward
- Format as Extended Journaled from this utility





At first the devices are JBOD
(Just a Bunch Of Drives)


1 Choose RAID


☐ Enhanced JBOD
☐ RAID Level 0
☐ RAID Level 1
☐ RAID Level 0+1
☒ RAID Level 5


5 RAID 5 offers the best mix of performance and data protection. One drive can fail without data loss.


Three or more drives are needed for this RAID level.

2 Select Drives


Bay 1
2 TB


Bay 2
2 TB



Bay 3
2 TB


Bay 4
2 TB

The total size of the RAID 5 set using the selected drives will be approximately 5.23 TB.


3 Options

☐ Use unassigned drives as spares

 Cancel Create

Select the Drives and Choose a RAID


I chose RAID Level 5 and all four drives (no spare)




The Configured RAID Set


Still need to create a volume

CONTROLLER


 **Status**
Initializing ●

 **Tasks**
1 task


VOLUMES


 **R1V1**
5.23 TB ●


RAID SETS


 **RS1**
5.23 TB ●

DRIVES

 **Bay 1**
2 TB ●

 **Bay 2**
2 TB ●

 **Bay 3**
2 TB ●

 **Bay 4**
2 TB ●

Task Initialize
Target Volume "R1V1" of RAID set "RS1"
Status Running

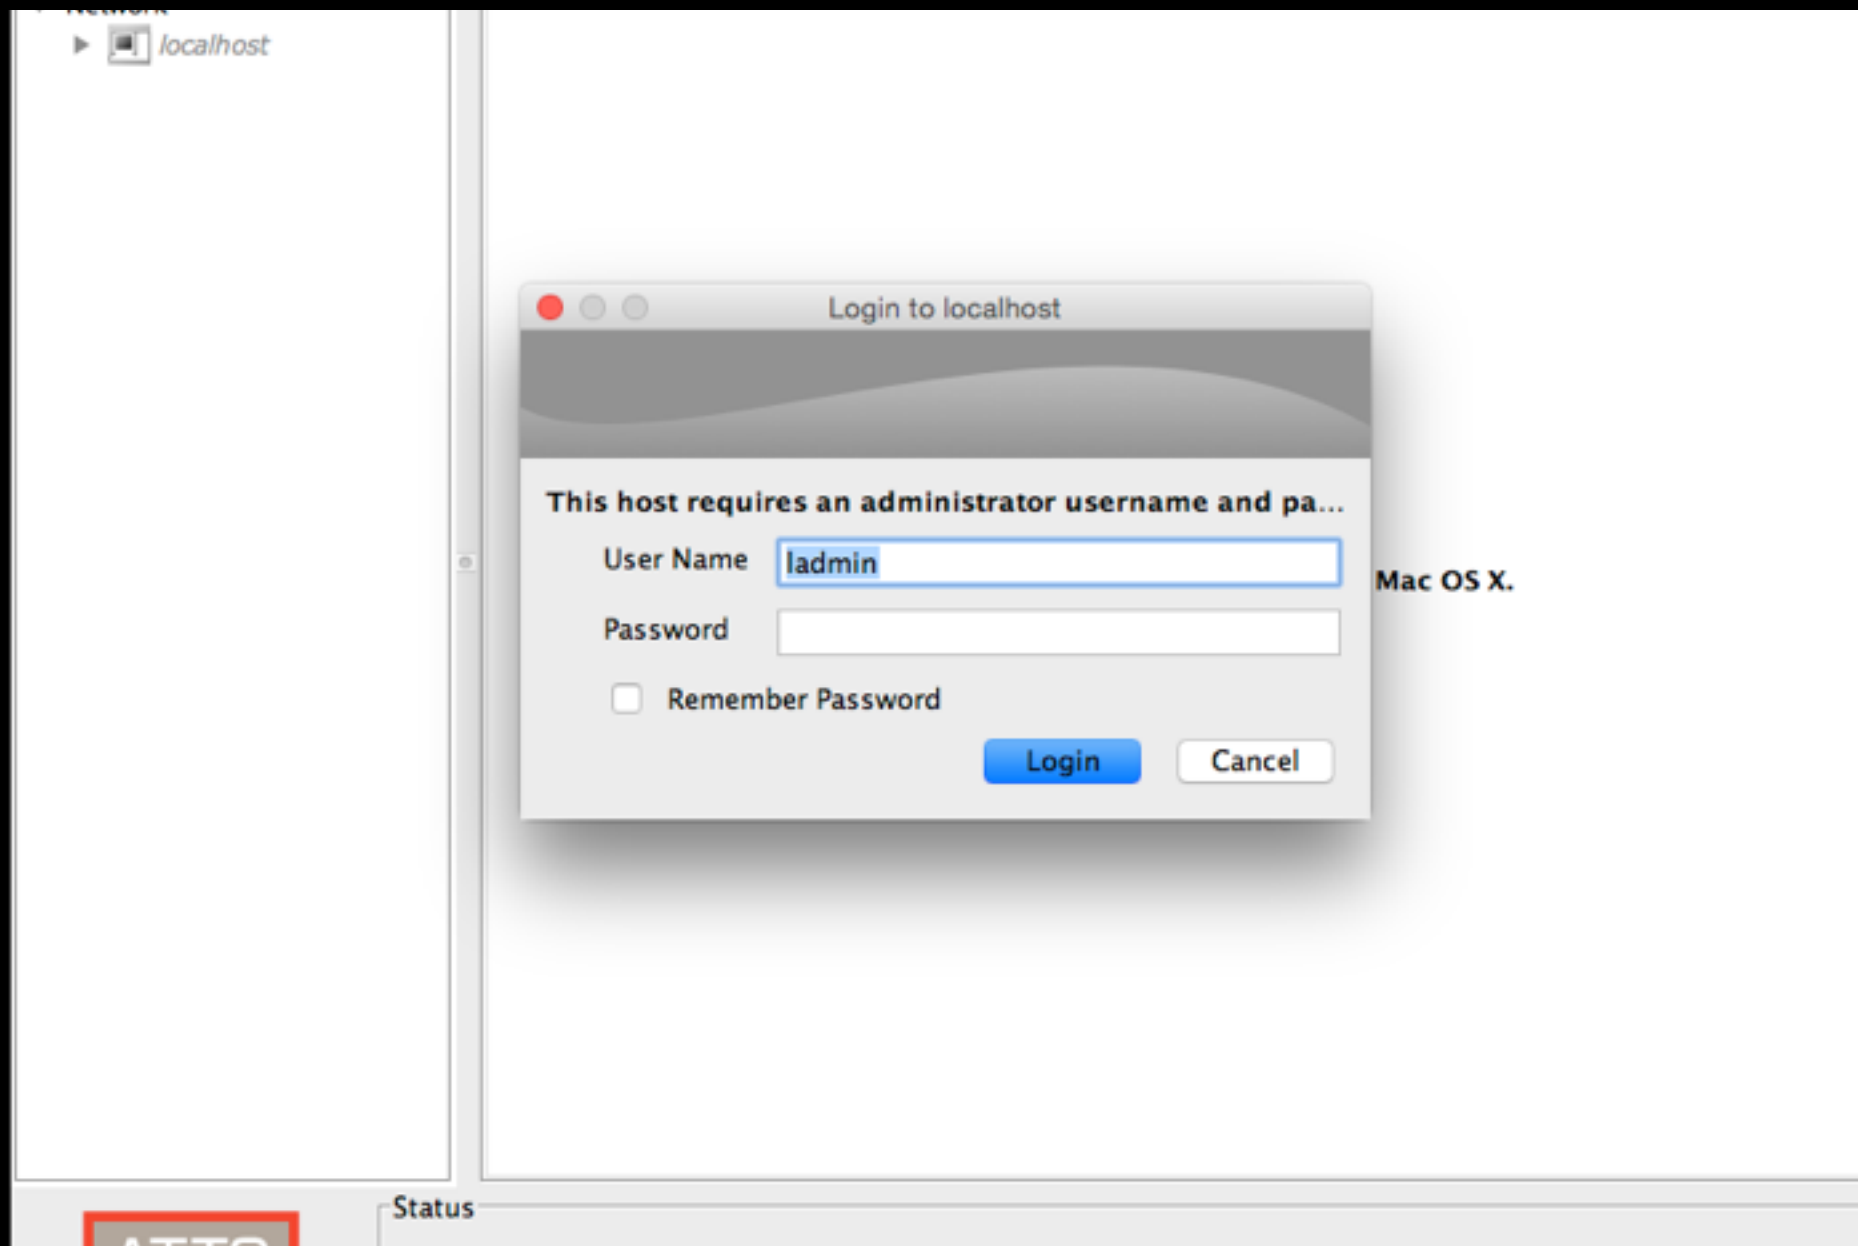
Creating a Volume

Monitor progress in Tasks

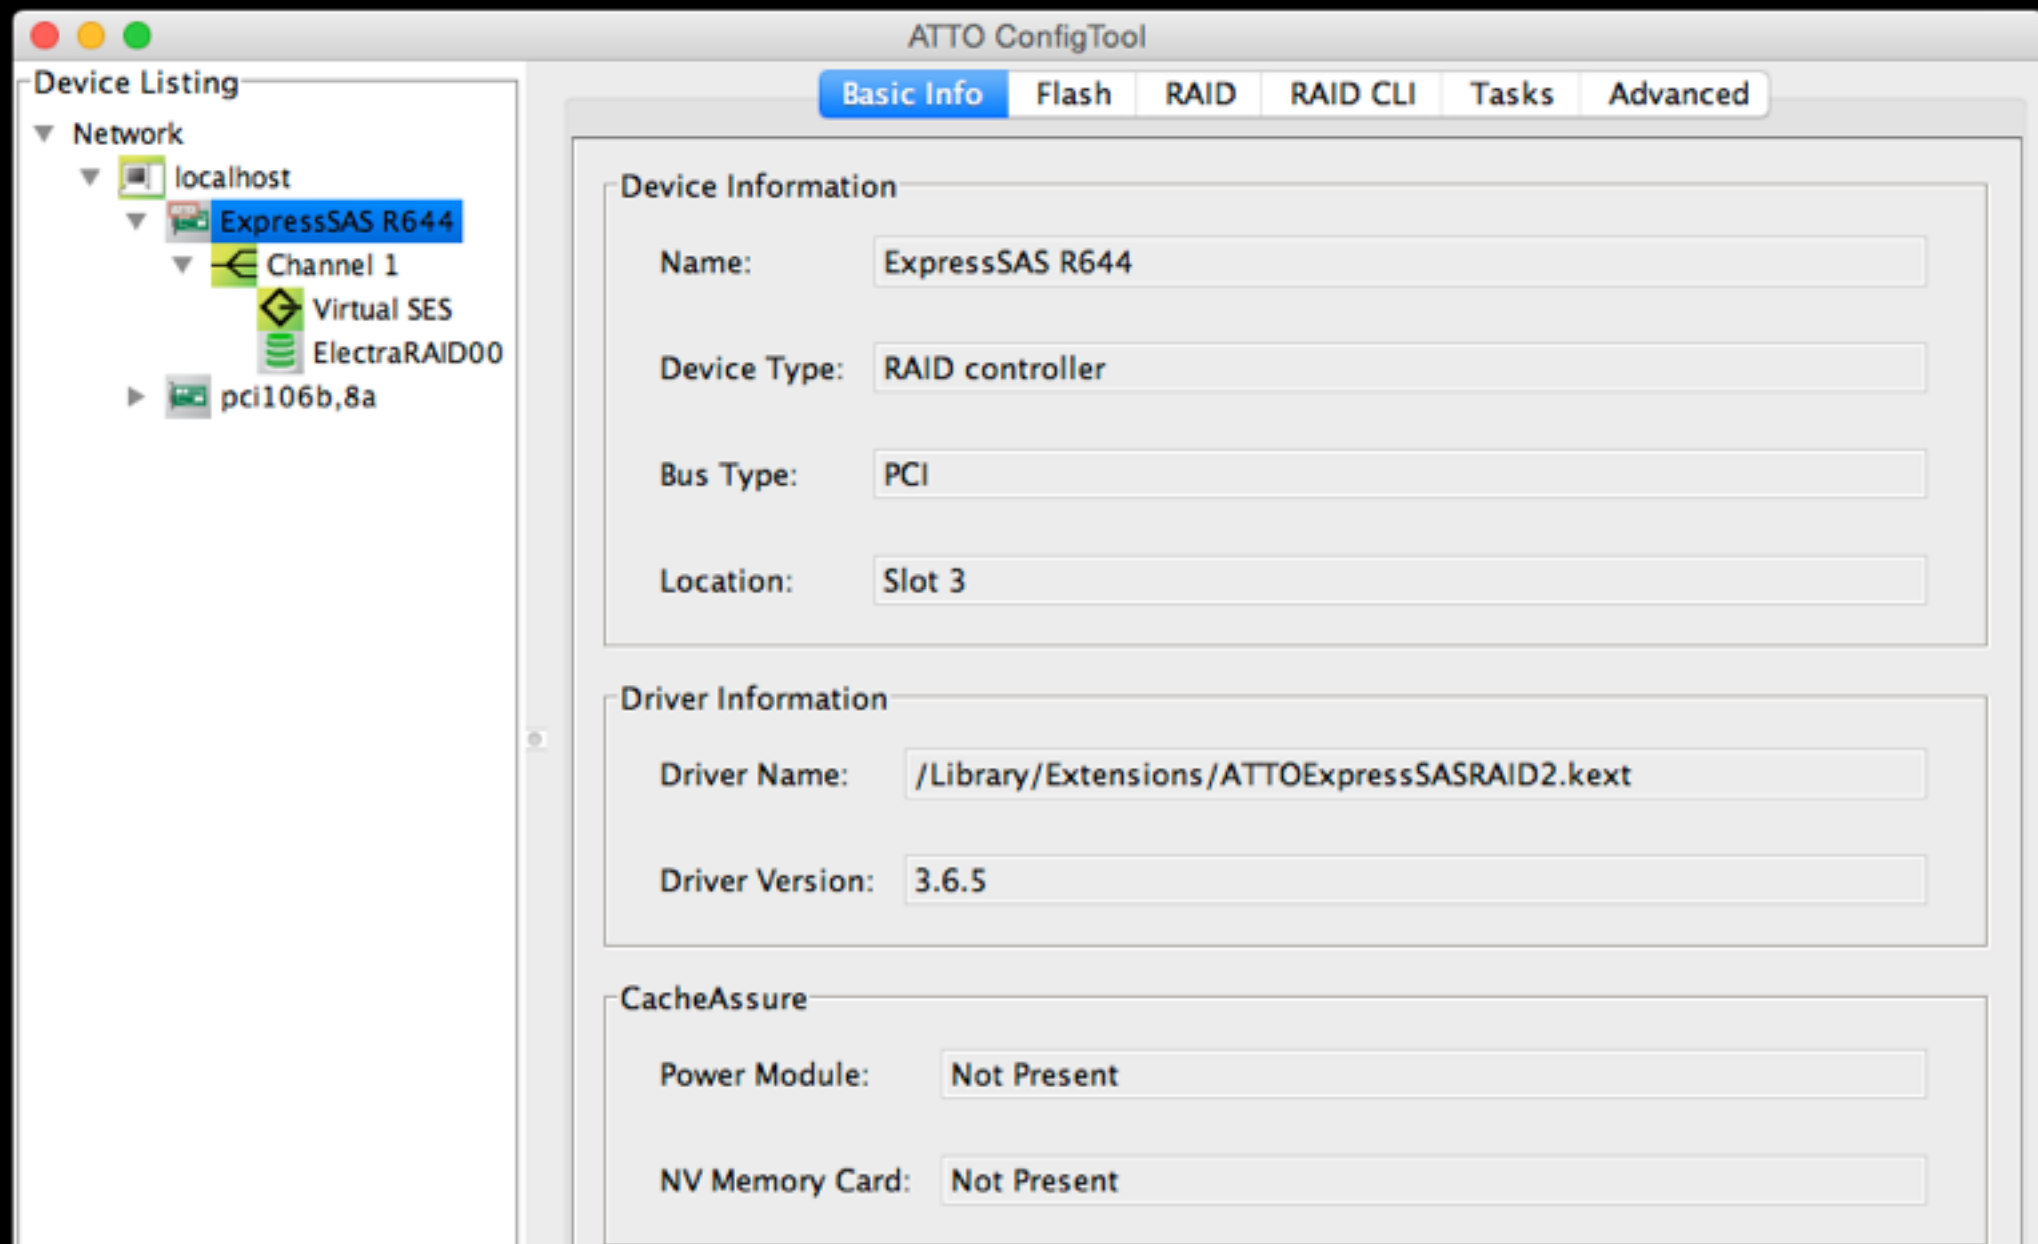
Create the SSD RAID

- Install the ATTO driver
- ATTO Config Tool requires Java 6 but is easy to use



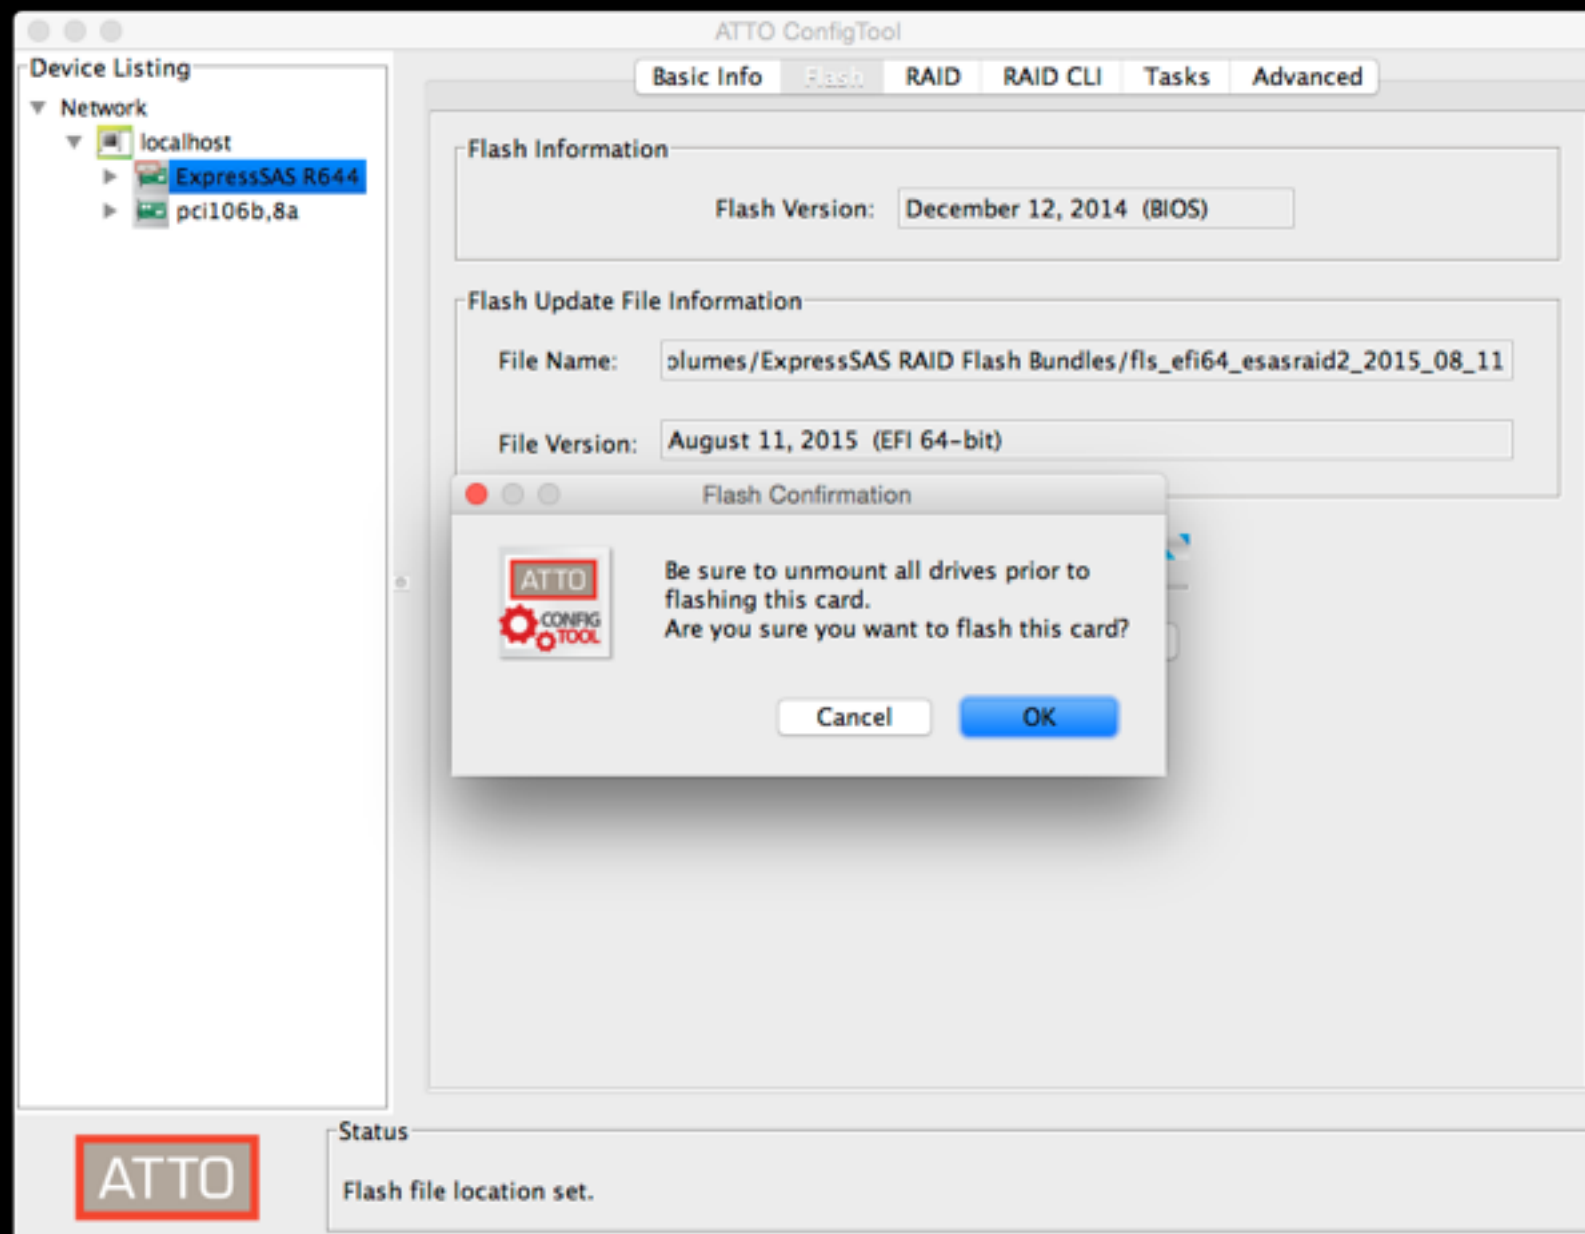


ATTO ConfigTool Requires Admin Access



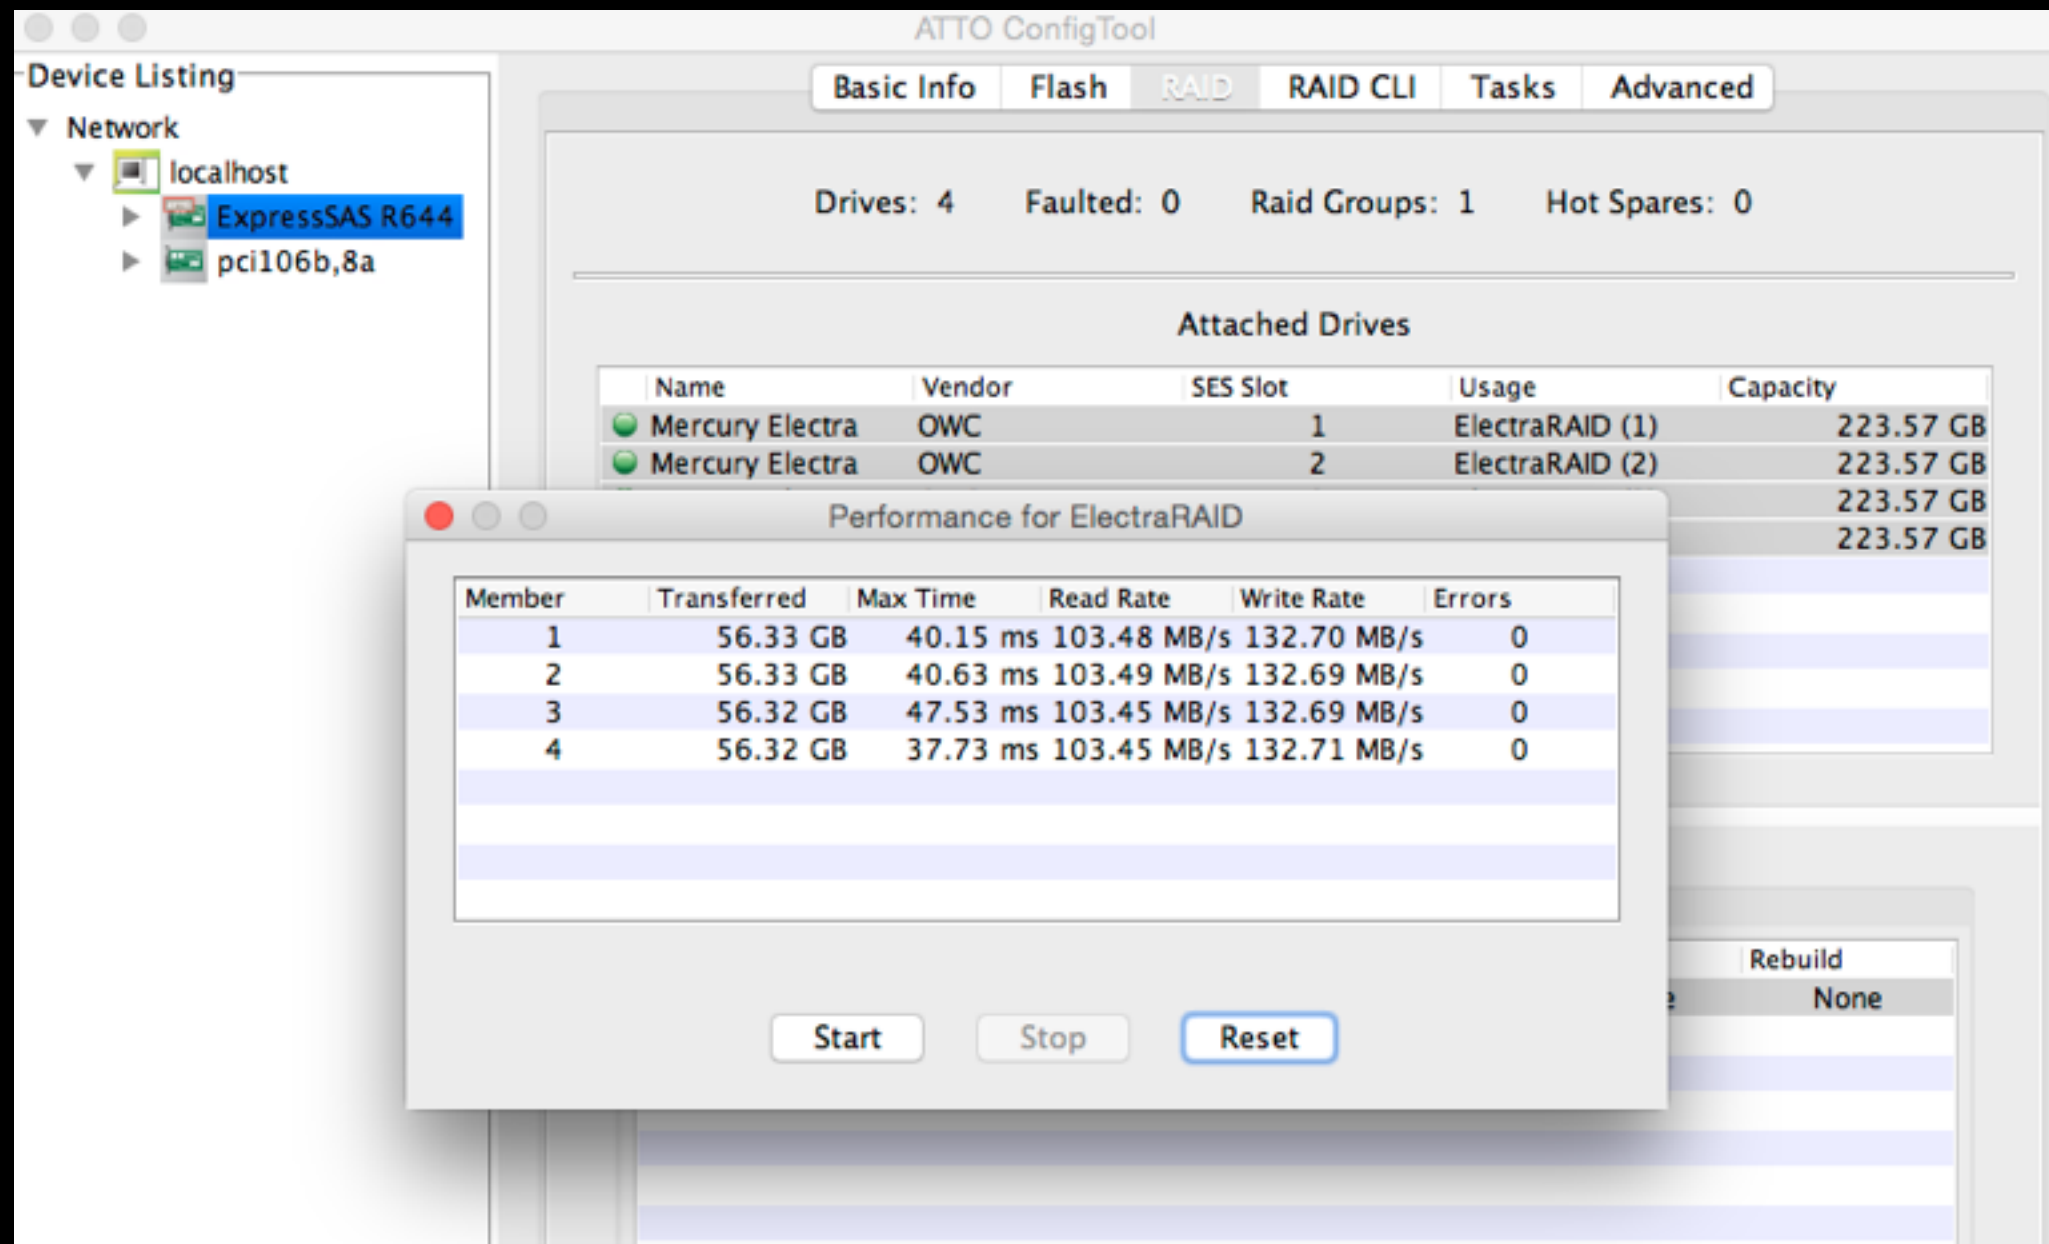
Verify Device and Driver

More on this later



Update RAID Controller Firmware

There are different versions of firmware for different computer systems.



Performance Test of RAID

I wasn't so much interested in the speed results as I just wanted an indication that the RAID was working at this point.

I had 2 independent
RAIDs, now what?

Install OS and Test Each RAID Independently

- Installed OS X Yosemite (and later El Capitan)
- Measured Boot Time with a Stopwatch
- Measure Data Speed with Blackmagic Disk Speed Test.app
- Results presented later

This is the best part.

Create the Fusion RAID Drive

- Three diskutil commands
 - list
 - corestorage create
 - corestorage createvolume

- diskutil list - use this to get the device nodes

Identify Devices

/dev/disk1

#:	TYPE	NAME	SIZE	
IDENTIFIER				
0:	GUID_partition_scheme		*5.2 TB	disk1
1:	EFI	EFI	209.7 MB	disk1s1
2:	Apple_HFS	Internal RAID temp	5.2 TB	disk1s2

/dev/disk5

#:	TYPE	NAME	SIZE	
IDENTIFIER				
0:	GUID_partition_scheme		*576.0 GB	disk5
1:	EFI	EFI	314.6 MB	disk5s1
2:	Apple_HFS	FastSSDRAID	575.6 GB	disk5s2

- diskutil corestorage create BigFastLVG /dev/
disk5 /dev/disk1

Identify the UUID

[lots of output omitted]

Waiting for Logical Volume Group to appear

Discovered new Logical Volume Group

"EF96CCA9-C7D7-4F03-BDF0-99C27D412D46"

Core Storage LVG UUID: EF96CCA9-C7D7-4F03-BDF0-99C27D412D46

Finished CoreStorage operation

- diskutil corestorage createvolume EF96CCA9-C7D7-4F03-BDF0-99C27D412D46 jhfs+ "Internal Storage" 100%

Identify the New Device

/dev/disk1

#:	TYPE	NAME	SIZE	IDENTIFIER
0:	GUID_partition_scheme		*5.2 TB	disk1
1:	EFI	EFI	209.7 MB	disk1s1
2:	Apple_CoreStorage		5.2 TB	disk1s2
3:	Apple_Boot	Boot OS X	134.2 MB	disk1s3

/dev/disk2

#:	TYPE	NAME	SIZE	IDENTIFIER
0:	Apple_HFS	Internal Storage	*5.8 TB	disk2
		Logical Volume on disk5s2, disk1s2		
		8EF44542-889E-414F-A7BA-1AC18F71C7DB		
		Unencrypted Fusion Drive		

/dev/disk5

#:	TYPE	NAME	SIZE	IDENTIFIER
0:	GUID_partition_scheme		*576.0 GB	disk5
1:	EFI	EFI	314.6 MB	disk5s1
2:	Apple_CoreStorage		575.6 GB	disk5s2
3:	Apple_Boot	Boot OS X	134.2 MB	disk5s3

Other Useful Commands

- `diskutil cs list`
- `diskutil cs info`

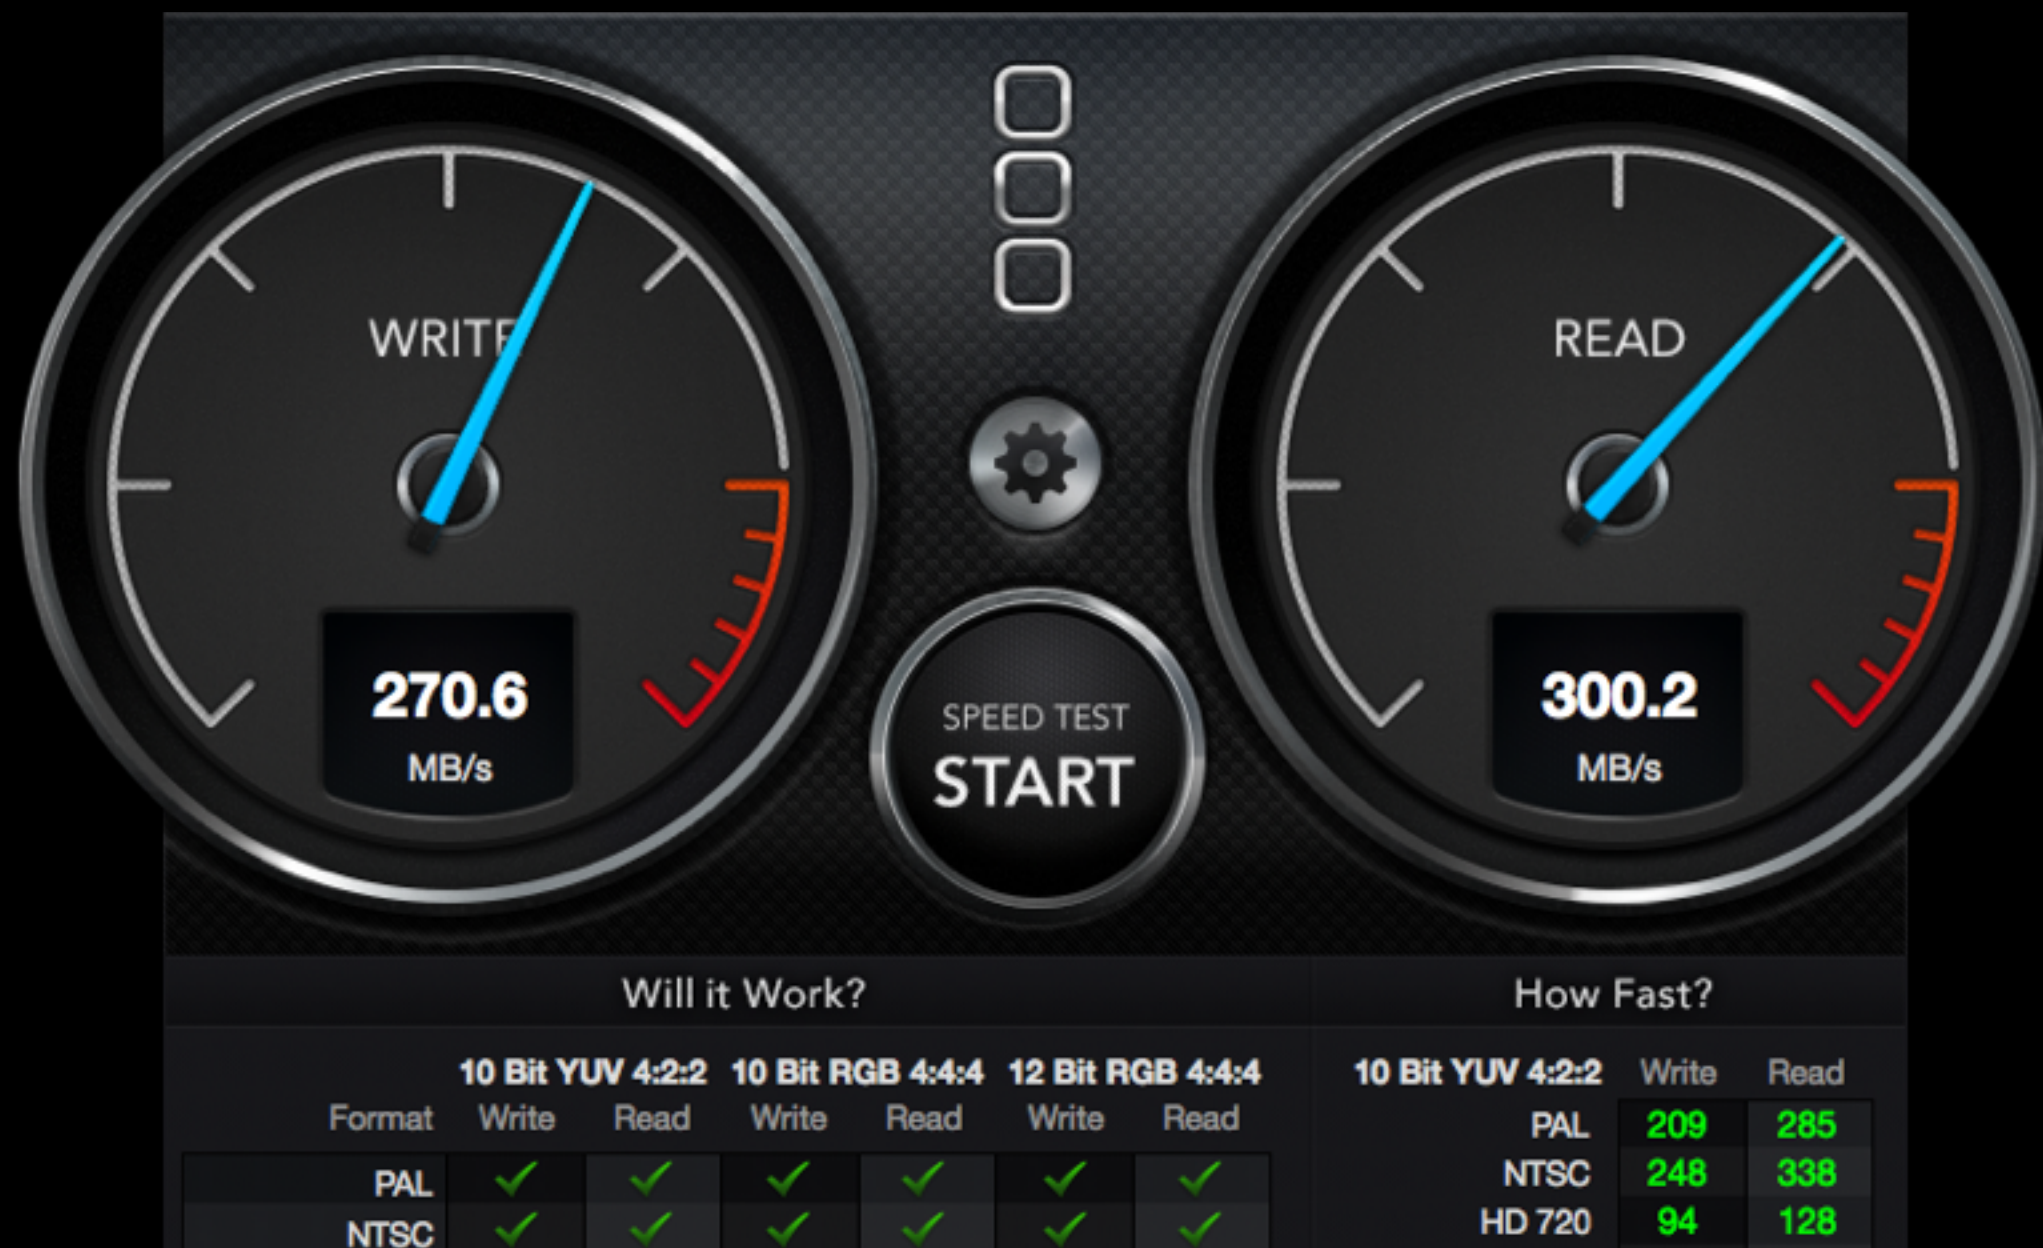
Repeat Installs...

- Had to reinstall...
 - OS X (of course)
 - Booted from an external FireWire drive
 - The install was relatively fast
- ATTO driver (kernel extension)
 - Install right after setup assistant before reboot

It's time to evaluate
the performance.


Test Results

	Apple Hardware RAID	ATTO SSD RAID	Fusion Drive made of RAIDs
Boot to Login Window	78 s	72 s	72 s
Storage Capacity	5.2 TB	.6 TB	5.8 TB




HD RAID only

Slowest in both read and write as expected



SSD RAID only

Fastest in read, but not writes oddly enough.



Fusion Drive of RAIDs

Fastest in writes, and close to SSD speed for reads.

Success!



Parts List


Part Description	Number	Cost per item	Extended Cost
240GB Mercury Electra 6G 2.5-inch 7mm SATA 6.0Gb/s Solid-State Drive	4	\$180	\$720
ATTO 644R RAID Controller	1	\$660	\$660
HITACHI 2TB 7200 R3.5 HD	4	\$50	\$200
ICY DOCK MB324SP-B ExpressCage 4x2.5" SAS/SATA HDD Hot Swap	1	\$45	\$45
1M Serial Attached SCSI SAS Cable SFF- 8087 4x Latching SATA	1	\$23	\$23
SATA 22 Pin Male to 7 Pin SATA Cable with 15 Pin SATA Female Power Cable	1	\$4	\$4
Apple RAID Card for MacPro	1	?	?

Tools List

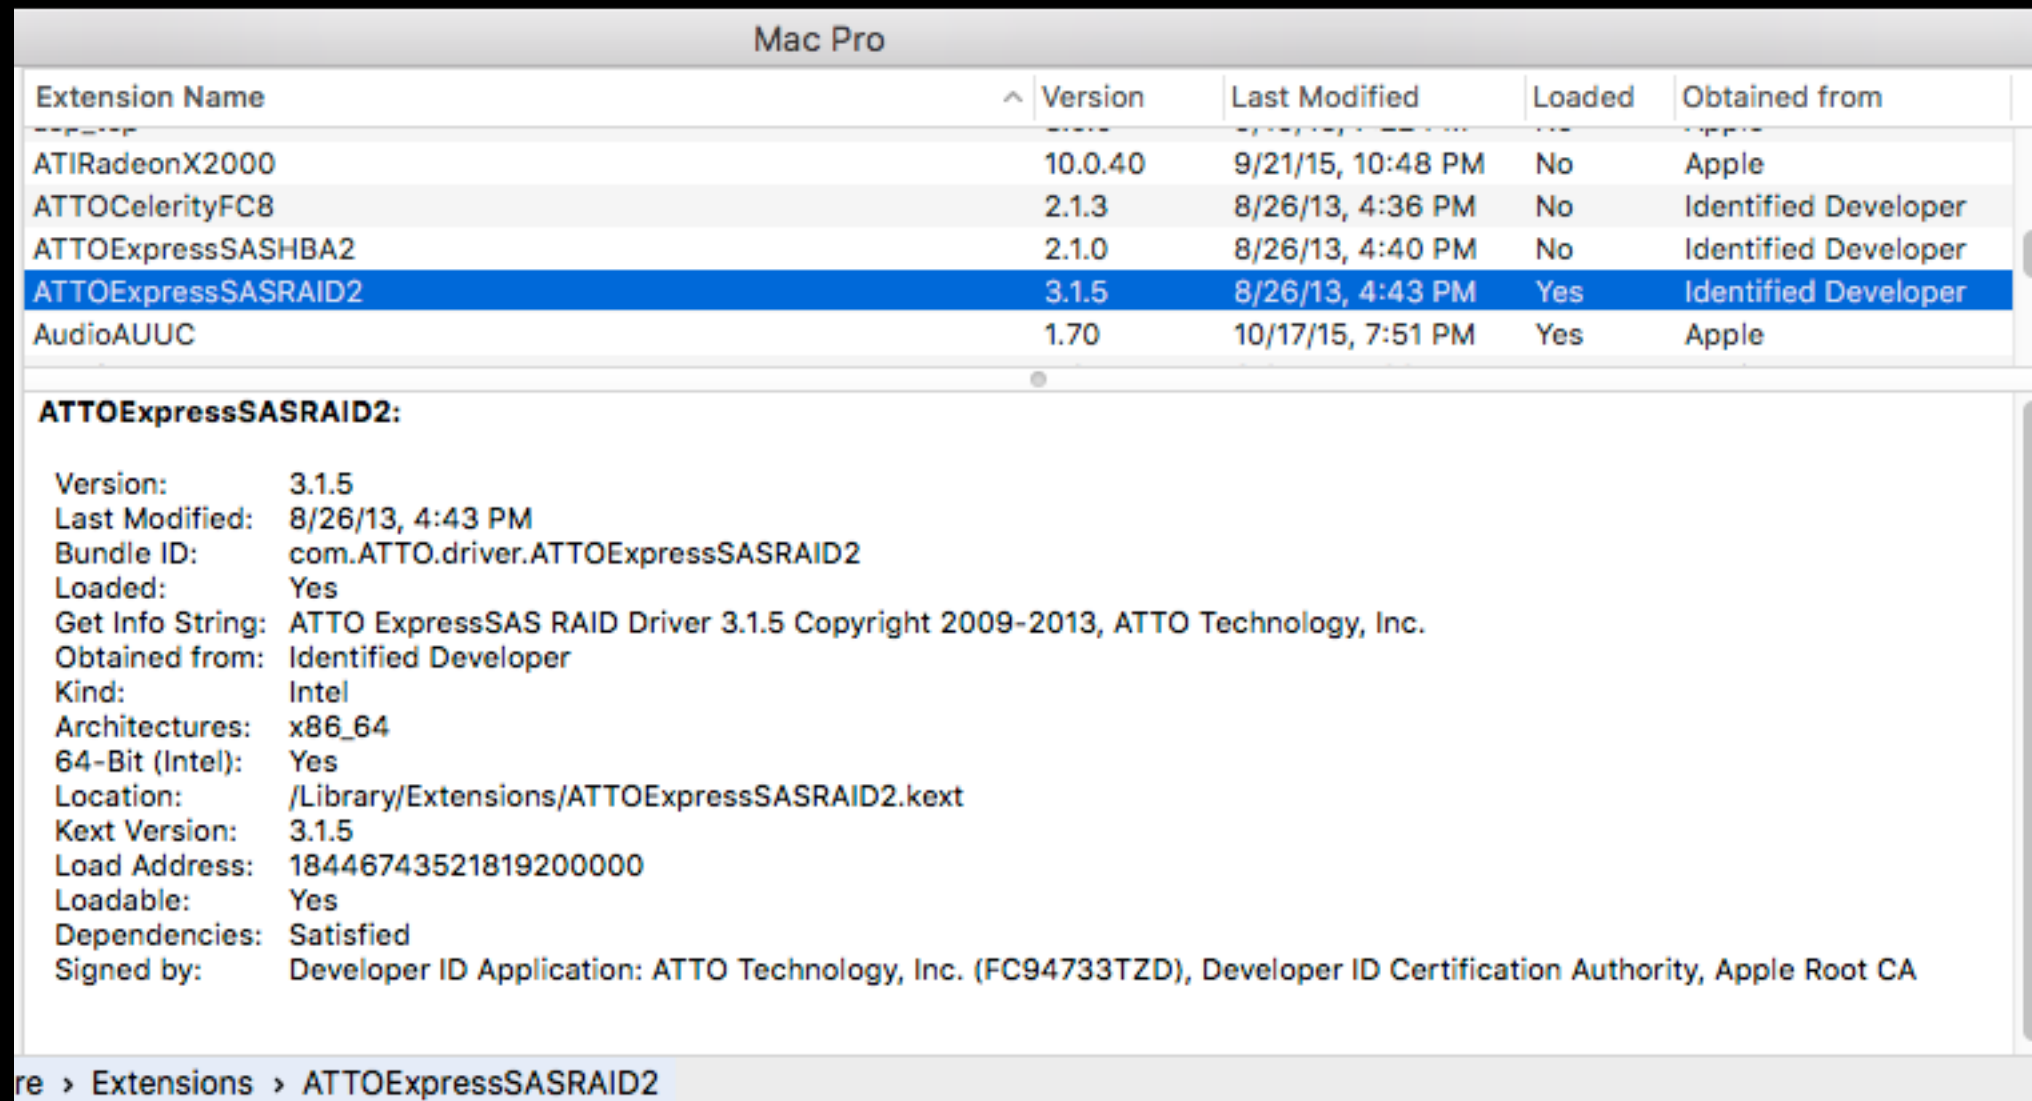
- #1 Phillips screw driver long
- #1 Phillips screw driver short

Lessons Learned

Icons Can be Confusing



Lessons Learned



Verify the Kernel Extension

According to ATTO documentation particular problems may be avoidable by editing the EFI firmware to change the timing of some parameters.

Lessons Learned

- Video editing with my new Fusion/RAID is much smoother than before.
- I still need more memory, maybe a faster GPU, etc, etc

Lessons Learned

- Cost savings of hardware upgrade vs purchase of a new computer system vary widely depending upon what you want to accomplish.
- I spent about \$1,700 on this project
- A new Mac Pro would cost at least \$2,999
- So I'm either \$1,700 behind on saving for my new dream system or I'm \$1,299 ahead.

Lessons Learned

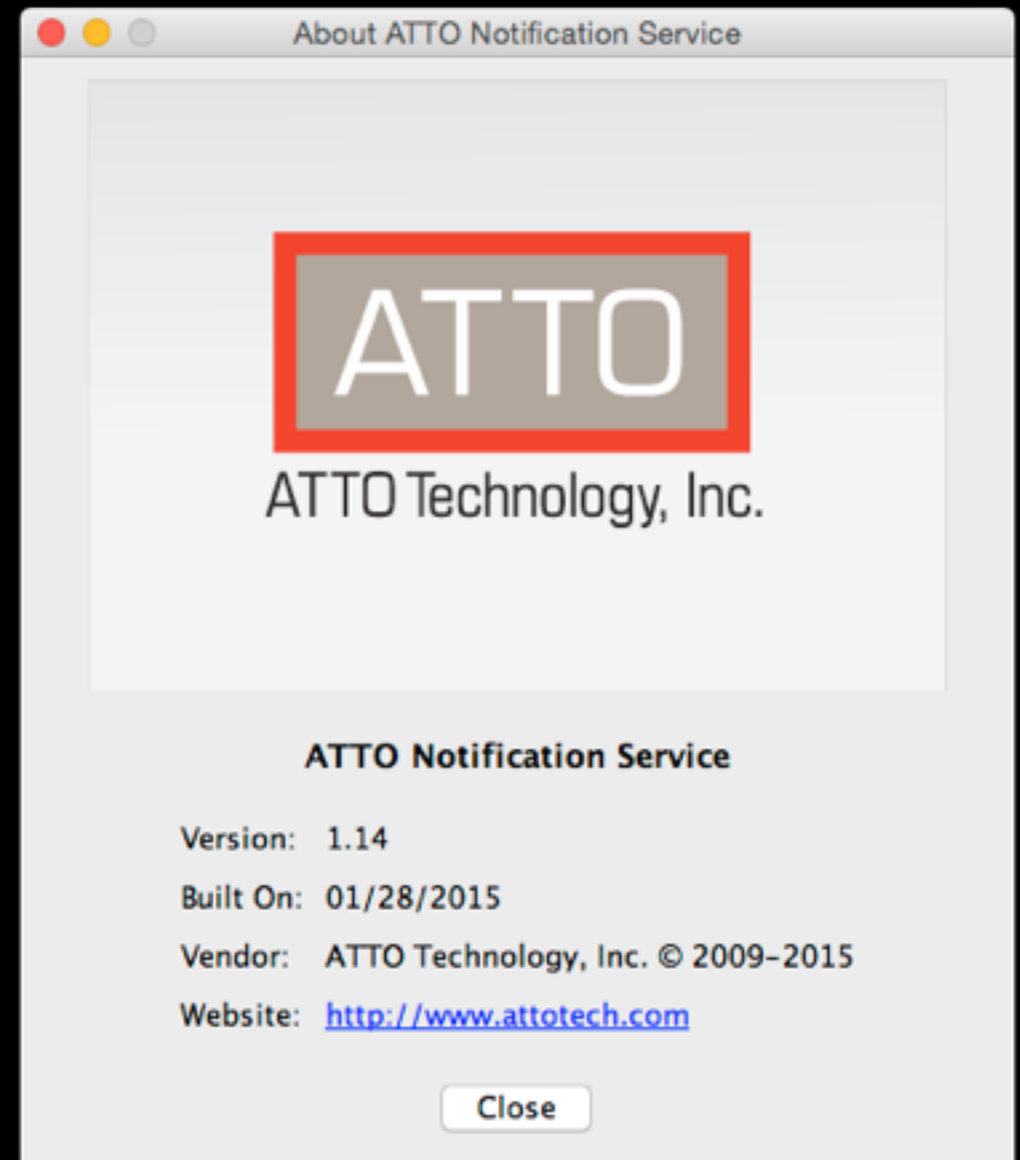
- Don't forget to increase backup storage size if increasing primary storage size.
- There is no substitute for a good backup.
- I have a external “Repair and Recovery” system drive with all necessary utilities and drivers ready to go if I encounter problems.
- And, did I mention the backup?

Lessons Learned

- He who hesitates will probably get a better, faster and cheaper computer (later) but won't have as much fun (now).
- You will make mistakes, but if you learn from them you will be better for it.

Lessons Learned

- It's important to monitor the health of the RAIDs
 - Apple RAID utility does this automatically
 - ATTO installs a notification tool (but still requires Java 6)
 - Also, Disk Utility in El Capitan seems to verify and repair Fusion drives better.



Lessons Learned

- When the RAIDs were independently configured and not part of a Fusion drive, the startup manager would not appear when the computer was started up and the option key was pressed
- This does not happen with the Fusion/RAID
- Again I may need to follow the ATTO documentation and edit EFI

Lessons Learned

- Fusion drives really don't like to be “separated”

This Fusion Drive is missing a disk.

Would you like to rebuild this Fusion Drive?

Rebuilding will erase the following disks:

Where to go from
here?

Future Improvements

- More memory, just as soon as I can.
- OWC offers a CPU upgrade by swapping out the motherboard.
- Even more storage, the ATTO RAID Controller has an external mini-SAS connector that could be used with an external SAS/SATA enclosure.
- How about more tiers? Is there a benefit to more tiers of storage?



Thanks for coming to
my session!

Chris R Tarnowiecky
@tarnythetrainer
tarny@tarnyinc.com