

Data syncing / Data distribution media
Originated by Bill Fanning (URI) on Dec 6, 2011

From: William L. Fanning (URI) on Tue, 06 Dec 2011

All,

What media are ships without multibeam using to distribute data to scientists?

We generate 4 to 5 Gigabytes of data per week. That's about one DVD per week of cruise. As you might guess burning multiple sets of DVDs for a month long cruise isn't much fun so we have been giving science their data on USB thumb drives.

Has anyone found a make/model of USB thumb drive that doesn't come with software pre-installed? Anyone using or have opinions on the pros and cons of SD cards? Any opinions on other options for data sets less than 32GB?

Thanks,
Bill

William L. Fanning
R/V Endeavor Technical Services
URI Graduate School of Oceanography
Narragansett, RI 02882

Reply From: Toby Martin (OSU) on Tue, 6 Dec 2011

> What media are ships without multibeam using to distribute data to scientists?

We are currently using 250-500 GB 2.5 inch USB harddisks.

> Has anyone found a make/model of USB thumb drive that doesn't come
> with software pre-installed? Anyone using or have opinions on the
> pros and cons of SD cards? Any opinions on other options for data
> sets less than 32GB?

Last year we had USB thumb drives, the 16 GB drives worked well, we had some problems with the 32 GB drives. Daryl has the specs on the drive types.

One major question we kicked about was what filesystem to use:

- FAT32: old clunky, doesn't support permissions, but most everything can read and write it;
- ExFAT: even slightly old windows systems can't dealt with it;
- ext: Windows can't dealt with it without extras;

- HFS: Windows can't deal with it without extras.
Currently we are using FAT32.

Toby

Reply From: Robbie Laird (WHOI) on Wed, 07 Dec 2011

Hi Bill

> What media are ships without multibeam using to distribute data to
> scientists?

We have multibeam, but how we do the data is unaffected by whether or not we are running it. (Although it's possible that our decision was driven by the largest data set we might run across, which is likely to be the multibeam.) We are using Transcend USB hard drives, I think either 320G or 500G. The model is the StoreJet 25M, which is the model with the rubber sleeve. It's marketed as having some degree of shock resistance. (I think the sleeve is actually silicone, so pretty much nothing sticks to it. For labelling, I sometimes slide an index card between the cover and the drive.) So far, we have not had any problems with the drives. We do ask that the drives be returned, but without any real expectation that they will always do so. Generally we get them back from R2R, and from WHOI archives, and some others.

On a multileg trip, you can put all the data on the one drive that is headed to archives. You can also set up the drive to mirror the data at some interval, this gives you (another?) backup, but also insures that the data is ready to deliver as soon as you hit port. (I suppose that applies to other mounted systems as well.)

The hard drives come with crap on the drive as well, we just ignore it. Eventually it gets erased when the drives comes back and gets erased in prep for re-use. As someone else mentioned, FAT filesystem seems to be the only way to go.

Robbie

Reply From: Thomas Wilson (SUNY) on Wed, 7 Dec 2011

Hi All,

I can second Robbie's endorsement of the Transcend Storejet drives. When my son Sean was a freshman in college I gave him one of these drives as a backup. Four years later he's run through two laptops and two desktop computers but still has and uses his original StoreJet.

I've recommended these to pretty much every parent sending a student to college for several years now

and have gotten no negative feedback. I'm certain it's possible to break this drive - but it isn't easy.

Regarding identification, although hardly anything sticks to the silicone case, there is a little hard plastic rim at one end of the drive that is just large enough for a 1/2 inch Brother P-Touch label.

Tom

Reply From: Jon C. Meyer (SIO) on December 6, 2011

>>> What media are ships without multibeam using to distribute data to scientists?

We are currently using 250-500 GB 2.5 inch USB harddisks.

We do the same, sometimes up to 1TB 2.5 inch HDDs for long/data intensive cruises -- they're only \$100 these days. However, during our pre-cruise meeting, we typically ask that the Science party bring their own HDD for this purpose. They usually are happy to do this; we have a stash on hand for when they don't bring their own.

One major question we kicked about was what filesystem to use:

- FAT32: old clunky, doesn't support permissions, but most everything can read and write it;
- ExFAT: even slightly old windows systems can't deal with it;
- ext: Windows can't deal with it without extras;
- HFS: Windows can't deal with it without extras.

Currently we are using FAT32.

> FAT32 negates rsync's usefulness in staging a lot of data, then doing a short sync at the end of the cruise.

The trick with using rsync on windows is to open up modify-window, since windows only has a 2 second tick in the timestamp.

```
rsync -rt --modify-window=3D2 source destination
```

does the trick nicely.

> So, we typically find out what the Chief Sci uses (usually NTFS, HFS or in rare cases ext3 or ext4), format as needed, and then uses a compatible system to copy the data from our server(s) to the HDD. We have Cygwin on our Windows systems, and any UNIX-like OS has rsync.

Reply from From: Thomas Wilson (SUNY) on Wed, 7 Dec 2011

For copying and syncing I recommend a program called Beyond Compare from scootersoftware.com. I've only used the Windows version but it is available for Linux as well.

GUI based, easy to learn the basics but extremely powerful. Allows various types of comparisons (name only, size, timestamp within x seconds) or full byte by byte comparison. Various options for copying, synchronizing, updating directories and directory trees including timestamp touch of otherwise identical files.

Handles huge jobs like a 100,000+ file archive transfer with ease. Very easy to pick up where you left off if for example a network or hardware glitch interrupts something halfway through. Fully functional trial version expires after you have used it on 30 different days by which time you will most likely want to pay the \$30 license fee.

Tom

Reply From: John Haverlack (UAF) on Wed, 7 Dec 2011

Hi all,

Other than the standard rsync tool which is available for linux, os x and windows I some times find DirSyncPro to be useful.

<http://www.dirsyncpro.org/>

It is:

- * Open Source
- * Free
- * Cross Platform (Win, Lin, Mac)
- * Has a user friendly GUI
- * Can also be scripted to automated synchronization
- * Fully functional for bi-directional or uni-directional synchronization

\$.02

Once nice feature of rsync is that it can throttle bandwidth which is nice if transferring large files over a shared slow connection.

--

John Haverlack
IT Manager, School of Fisheries and Ocean Sciences
University of Alaska Fairbanks
Fairbanks, Alaska 99775-7220

Reply From: Webb Pinner (URI) on Wed, 7 Dec 2011

Another rsync reference i wrote awhile back for setting up rsync jobs as scheduled tasks

<http://www.oceandatarat.org/?p=29>

- Webb

Reply From: Kurt Schwehr (UNH) on Wed, 7 Dec 2011

What Webb said. Has anyone used these two options with success?

--partial --append-verify

and don't forget

--bwlimit=KBPS

BTW, As noted by Dan Christian, one of my fellow sat data delivery protocol co-designers:

"High latency and drop outs drive TCP crazy. It thinks every dropped packet means congestion and backs off. Also, you need to increase the window size."

My take is that the real solution really should be something like bittorrent... but with UDP and an understanding of how long/short and fat/skinny each link is and a preference to only send out packets to one destination for the sat links.

http://schwehr.org/blog/archives/2011-09.html#e2011-09-30T07_35_03.txt

-kurt

Reply From: Webb Pinner (URI) on Wed, 7 Dec 2011

The use of partial (-P) over satellite links is clutch. I also highly recommend the -z flag which uses the gzip compression algorithm on all out going traffic. For ASCII data this results on ~8x the data throughput.

According to Steve from hi-seas, the modems now support TCP proxy which essentially converts TCP traffic into UDP streams. This helps reduce the bottlenecks caused by tcp ack bits.

-W

Reply From: David Gassier (LDEO) on Wed, 7 Dec 2011

Hi,

I've been using this tool intensively for very important data control and verification after any data transfer:

<http://www.quickpar.org.uk/>

When you have a lot of data, this software generates files to repair your data set in case they are corrupted after a transfer.

You run the software on the dataset, it creates files of about 10% the size of your data (you actually set the data recovery trigger).

So when you transmit your dataset, you also transmit these files. These extra files will allow the scientist/the end user to verify that the data is not corrupted due to the media your using.

These files take some time to be generated but it is worth a try especially with the type of data we are talking about.

If you have these PAR2 files on a corrupted disk drive you'd be able to fix and repair them.

Give it a try it's very powerful. The first time I tried it I thought it was magic...

David

Reply From: Val Schmidt (UNH) on Thu, 8 Dec 2011

Kurt, what is it about bittorrent would help here? My (perhaps mis-) understanding of bittorrent is that it requires many sources from which to pull.

Also to add to the general discussion, I've read and found myself that operating many rsync sessions simultaneously is faster than a single transfer on the same data set. It's also faster to tar up the data before transfer if you have loads and loads of data files. For example, we might collect 10's of 1000's of photographs on an AUV mission. It's faster to tar up the directory tree and transfer than point rsync at the top and execute recursively.

-Val

Reply From: Kurt Schwehr (UNH) on Thu, 8 Dec 2011

Val,

Good question ▣ I look forward to the day when data is never again shipped more than a couple hundred feet by DVD/flashdrive/USB Drive/Tape (or maybe a helo flight from a ship in the ice to

McMurdo). Our data is not a Washington Mutual data center armored truck sized data center migration (a story I very much enjoyed hearing).

My thoughts so far...

Benefits of a bittorrent style setup:

- designed with networks that frequently go up and down. rsync over ssh will work, but not great
- It's okay with transferring bits and pieces
- If you hit a high bandwidth spot like wifi or wired in a harbor, it will happily ramp up the bandwidth.
- Think Barrow, AK... it could do a fast dump to a local dump to a shore client when a ship came in range and now you have two seeds getting data back through the system to the national archives
- This can be hands off... the ship just sends the file definition / seed info over the sat link and then the archive sites can monitor the progress and the ship slowly and sporadically sends the files in no particular order.
- If two ships passed using SWAP, it is possible to pass data to a vessel that will hit higher bandwidth internet sooner.
- Again with SWAP, if you have one ship with a much bigger sat pipe, bittorrent could totally handle the UUCP like job of bouncing the data to shore.

Downsides:

- What you don't want (that is in the current spec) is to have more than 1 or 2 clients pulling that data. Having clients "swarm" the ships network needs to be prevented.
- Most IT groups will so "no way" to anything peer-to-peer without ever looking beyond the first time they see a word proposing such.

What needs to happen? IT teams need to be willing to consider a peer-to-peer system without a knee jerk no. I do understand that plain old bittorrent has a well deserved bad rap for 95% of its traffic. Stock bittorrent setup would definitely work and be able to be setup to stay away for the normal bittorrent world. But it would be better to have a modified bittorrent setup that:

- Used UDP to send and I'm not sure if it should be TCP or UDP to "ack" blocks
- Have a configurable or dynamic UDP packet size to best work with the specific satellite configuration and bit error rate
- Can listen to the ship's systems for hints about how much sat bandwidth to try to use
- Being able to know how many clients to send data to based on the route to that client (e.g. over sat verses land/sea based WAN)
- Evaluate if multicast could reduce the number of retransmits needed by increasing the chance that someone would get the packet and ack it
- Automate the seek info handling and distribution
- Can we use this channel to more effectively update ship board software systems? e.g. data delivery to ships
- A priority system to make sure that the most important data files or data types get to shore first (and I would imagine that will always be changing based on the vessel science goals)

The other thought I should bring up... I strongly suggest pre-compressing all data files before doing

bittorrent / rsync / whatever. Most of the formats compress very well, but that adds time to the transmit process (more blocks for bittorrent to track or more data for rsync to checksum during each time rsync runs). Adding a "nohup bzip2 -9 monster-mb-data-file.all" to your processing chain.

I'd be interested to hear what others thought about this (good/bad/otherwise ☐). I don't have available engineering time to work on such a system, but it would be great to flush out the idea and then we can have it out there with the chance that an enterprising Computer Science grad student in networking might take it on as a thesis project.

-kurt

Reply From: Steve Foley (SIO) on December 9, 2011

The accelerators do something like this.

For that matter, they also compress data on the fly, and will packet cache streams that aren't already compressed. I haven't run the numbers in reality, but it may actually be that you get better performance if you DONT compress at the rsync level. In reality, my guess is that the gzip algorithm is a little better than what gets used by the accelerator, and any caching that would have happened in a non compressed stream is already compressed out. So...dont stress not tossing in the -z if you have your accelerator going.

-Steve