

Article development led by [acmqueue](http://queue.acm.org)
queue.acm.org

**The accepted wisdom
does not always hold true.**

BY SACHIN DATE

Should You Upload or Ship Big Data to the Cloud?

IT IS ACCEPTED wisdom that when the data you wish to move into the cloud is at terabyte scale and beyond, you are better off shipping it to the cloud provider, rather than uploading it. This article takes an analytical look at how shipping and uploading strategies compare, the various factors on which they depend, and under what circumstances you are better off shipping rather than uploading data, and vice versa. Such an analytical determination is important to make, given the increasing availability of gigabit-speed Internet connections, along with the explosive growth in data-transfer speeds supported by newer editions of drive interfaces such as SAS and PCI Express. As this article reveals, the aforementioned “accepted wisdom” does not always hold true, and there are well-reasoned, practical recommendations

for uploading versus shipping data to the cloud.

Here are a few key insights to consider when deciding whether to upload or ship:

- ▶ A direct upload of big data to the cloud can require an unacceptable amount of time, even over Internet connections of 100Mbps (megabits per second) and faster. A convenient workaround has been to copy the data to storage tapes or hard drives and ship it to the cloud datacenter.

- ▶ With the increasing availability of affordable, optical fiber-based Internet connections, however, shipping the data via drives becomes quickly unattractive from the point of view of both cost and speed of transfer.

- ▶ Shipping big data is realistic only if you can copy the data into (and out of) the storage appliance at very high speeds and you have a high-capacity, reusable storage appliance at your disposal. In this case, the shipping strategy can easily beat even optical fiber-based data upload on speed, provided the size of data is above a certain threshold value.

- ▶ For a given value of drive-to-drive data-transfer speed, this threshold data size (beyond which shipping the data to the cloud becomes faster than uploading it) grows with every Mbps increase in the available upload speed. This growth continues up to a certain threshold upload speed. If your ISP provides an upload speed of greater or equal to this threshold speed, uploading the data will always be faster than shipping it to the cloud, no matter how big the data is.

Suppose you want to upload your video collection into the public cloud; or let's say your company wishes to migrate its data from a private datacenter to a public cloud, or move it from one datacenter to another. In a way it doesn't matter what your profile is. Given the explosion in the amount of digital information that both individuals and enterprises have to deal with, the prospect of moving big data from one place to another



over the Internet is closer than you might think.

To illustrate, let's say you have 1TB of business data to migrate to cloud storage from your self-managed datacenter. You are signed up with a business plan with your ISP that guarantees you an upload speed of 50Mbps and a download speed of 10 times as much. All you need to do is announce a short system-down-time window and begin hauling your data up to the cloud. Right?

Not quite.

For starters, you will need a whopping 47 hours to finish uploading 1TB of data at a speed of 50Mbps—and that's assuming your connection never drops or slows down.

If you upgrade to a faster—say, 100Mbps—upload plan, you can finish

the job in one day. But what if you have 2 TB of content to upload, or 4TB, or 10TB? Even at a 100Mbps sustained data-transfer rate, you will need a mind-boggling 233 hours to move 10TB of content!

As you can see, conventional wisdom breaks down at terabyte and petabyte scales. It's necessary to look at alternative, nonobvious ways of dealing with data of this magnitude.

Here are two such alternatives available today for moving big data:

- ▶ Copy the data locally to a storage appliance such as LTO (linear tape open) tape, HDD (hard-disk drive), or SSD (solid-state drive), and ship it to the cloud provider. For convenience, let's call this strategy "Ship It!"

- ▶ Perform a cloud-to-cloud transfer

of content over the Internet using APIs (application programming interfaces available) from both the source and destination cloud providers.⁶ Let's call this strategy "Transfer It!"

This article compares these alternatives, with respect to time and cost, to the baseline technique of uploading the data to the cloud server using an Internet connection. This baseline technique is called "Upload It!" for short.

A Real-Life Scenario

Suppose you want to upload your content into, purely for the sake of illustration, the Amazon S3 (Simple Storage Service) cloud, specifically its datacenter in Oregon.² This could well be any other cloud-storage service provided by players⁹ in this space such as (but not limited to) Microsoft, Google, Rackspace, and IBM. Also, let's assume that your private datacenter is located in Kansas City, MO, which happens to be approximately geographically equidistant from Amazon's datacenters² located in the eastern and western U.S.

Kansas City is also one of the few places where a gigabit-speed optical-fiber service is available in the U.S. In this case, it's offered by Google Fiber.⁷

As of November 2015, Google Fiber offers one of the highest speeds that an ISP can provide in the U.S.: 1Gbps (gigabit per second), for both upload and download.¹³ Short of having access to a leased Gigabit Ethernet¹¹ line, an optical fiber-based Internet service is a really, really fast way to shove bits up and down Internet pipes anywhere in the world.

Assuming an average sustained upload speed of 800Mbps on such a fiber-based connection,¹³ (that is, 80% of its advertised theoretical maximum speed of 1Gbps), uploading 1TB of data will require almost three hours to upload from Kansas City to S3 storage in Oregon. This is actually pretty quick (assuming, of course, your connection never slows down). Moreover, as the size of the data increases, the upload time increases in the same ratio: 20TB requires 2.5 days to upload, 50TB requires almost a week to upload, and 100TB requires twice that long. At the other end of the scale, a half a petabyte of data requires two

Figure 1. Data flow when copying data to a storage appliance.

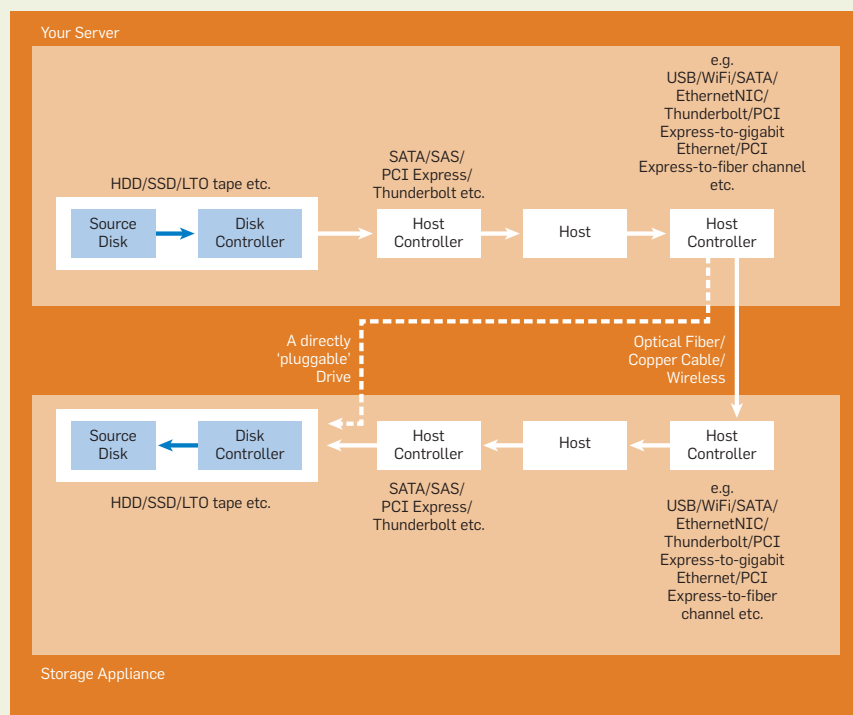


Figure 2. Data transfer speeds supported by various interfaces.

Interface Type	Data Transfer Speed (Gbps)
SATA Revision 3	6 ¹⁷
SAS-3	12 ¹⁰
SuperSpeed USB (USB 3.0)	10 ²⁰
PCI Express version 4	15.754 (single data lane), to 252.064 (16 data lanes) ¹⁴
Thunderbolt 2	20 ¹

months to upload. Uploading one petabyte at 800Mbps should keep you going for four months.

It's time to consider an alternative.

Ship It!

That alternative is copying the data to a storage appliance and shipping the appliance to the datacenter, at which end the data is copied to cloud storage. This is the Ship It! strategy. Under what circumstances is this a viable alternative to uploading the data directly into the cloud?

The mathematics of shipping data.

When data is read out from a drive, it travels from the physical drive hardware (for example, the HDD platter) to the on-board disk controller (the electronic circuitry on the drive). From there the data travels to the host controller (a.k.a. the host bus adapter, a.k.a. the interface card) and finally to the host system (for example, the computer with which the drive is interfaced). When data is written to the drive, it follows the reverse route.


When data is copied from a server to a storage appliance (or vice versa), the data must travel through an additional physical layer, such as an Ethernet or USB connection existing between the server and the storage appliance.

Figure 1 is a simplified view of the data flow when copying data to a storage appliance. The direction of data flow shown in the figure is conceptually reversed when the data is copied out from the storage appliance to the cloud server.


Note that often the storage appliance may be nothing more than a single hard drive, in which case the data flow from the server to this drive is basically along the dotted line in the figure.

Given this data flow, a simple way to express the time needed to transfer the data to the cloud using the Ship It! strategy is shown in Equation 1, where: V_{content} is the volume of data to be transferred in megabytes (MB).

$\text{Speed}_{\text{copyIn}}$ is the *sustained* rate in MBps (megabytes per second) at which data is copied from the source drives to the storage appliance. This speed is essentially the minimum of three speeds: the speed at which the controller reads data out of the source drive and transfers it to the



Given the explosion in the amount of digital information that both individuals and enterprises have to deal with, the prospect of moving big data from one place to another over the Internet is closer than you might think.



host computer with which it interfaces; the speed at which the storage appliance's controller receives data from its interfaced host and writes it into the storage appliance; and the speed of data transfer between the two hosts. For example, if the two hosts are connected over a Gigabit Ethernet or a Fibre Channel connection, and the storage appliance is capable of writing data at 600MBps, but if the source drive and its controller can emit data at only 20MBps, then the effective copy-in speed can be at most 20MBps.

$\text{Speed}_{\text{copyOut}}$ is similarly the *sustained* rate in MBps at which data is copied out of the storage appliance and written into cloud storage.

T_{transit} is the transit time for the shipment via the courier service from source to destination in hours.

T_{overhead} is the overhead time in hours. This can include the time required to buy the storage devices (for example, tapes), set them up for data transfer, pack and create the shipment, and drop it off at the shipper's location. At the receiving end, it includes the time needed to process the shipment received from the shipper, store it temporarily, unpack it, and set it up for data transfer.

The use of sustained data-transfer rates. Storage devices come in a variety of types such as HDD, SSD, and LTO. Each type is available in different configurations such as a RAID (redundant array of independent disks) of HDDs or SSDs, or an HDD-SSD combination where one or more SSDs are used as a fast read-ahead cache for the HDD array. There are also many different data-transfer interfaces such as SCSI (Small Computer System Interface), SATA (Serial AT Attachment), SAS (Serial Attached SCSI), USB, PCI (Peripheral Component Interconnect) Express, Thunderbolt, and so on. Each of these interfaces supports a different theoretical maximum data-transfer speed.

Figure 2 lists the data-transfer speeds supported by a recent edition of some of these controller interfaces.

The effective copy-in/copy-out speed while copying data to/from a storage appliance depends on a number of factors:

- *Type of drive.* For example, SSDs are usually faster than HDDs partly because of the absence of any mov-

ing parts. Among HDDs, higher-RPM drives can exhibit lower seek times | higher areal-density (bits per surface area) drives can lead to higher data-transfer rates.

Figure 3. Growth in data transfer time, 800Mbps vs. tapes.

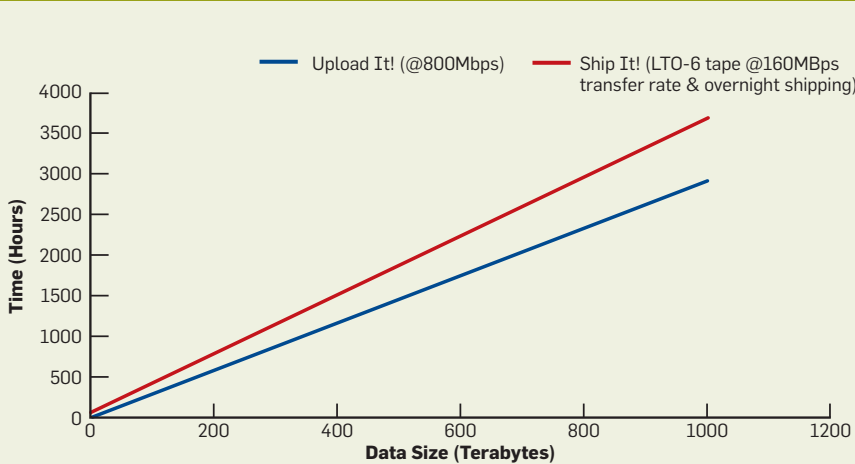


Figure 4. Growth in data transfer time, 100Mbps vs. tapes.


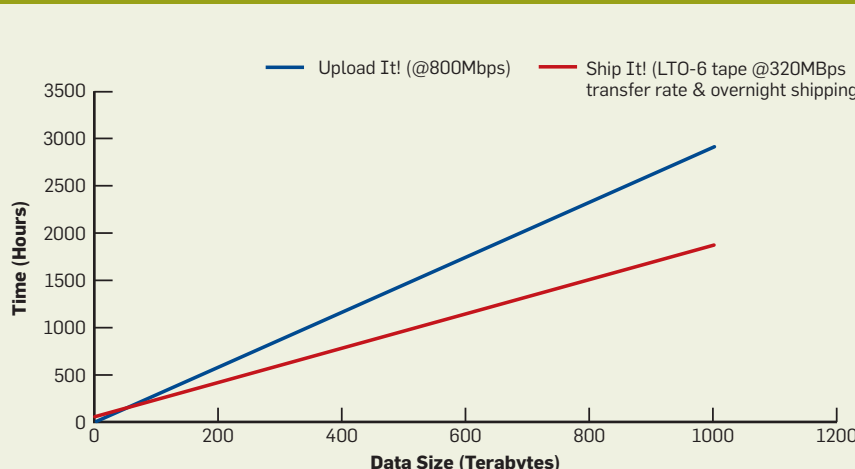


Figure 5. Growth in data transfer time, 800Mbps vs. 320MBps.



► *Configuration of the drive.* Speeds are affected by, for example, single disk versus an array of redundant disks, and the presence or absence of read-ahead caches on the drive.

► *Location of the data on the drive.* If the drive is fragmented (particularly applicable to HDDs), it can take longer to read data from and write data to it. Similarly, on HDD platters, data located near the periphery of the platter will be read faster than data located near the spindle. This is because the linear speed of the platter near the periphery is much higher than near the spindle.

► *Type of data-transfer interface.* SAS-3 versus SATA Revision 3, for example, can make a difference in speeds.

► *Compression and encryption.* Compression and/or encryption at source and decompression and/or de-encryption at the destination reduce the effective data-transfer rate.

Because of these factors, the effective *sustained* copy-in or copy-out rate is likely to be much different (usually much less) than the burst read/write rate of either the source drive and its interface or the destination drive and its controller interface.

With these considerations in mind, let's run some numbers through Equation 1, considering the following scenario. You decide to use LTO-6 tapes for copying data. An LTO-6 cartridge can store 2.5TB of data in uncompressed form.¹⁸ LTO-6 supports an uncompressed read/write data speed of 160MBps.¹⁹ Let's make an important simplifying assumption that both the source drive and the destination cloud storage can match the 160MBps transfer speed of the LTO-6 tape drive (that is, $\text{Speed}_{\text{copyIn}} = \text{Speed}_{\text{copyOut}} = 160 \text{ MBps}$). You choose the overnight shipping option and the shipper requires 16 hours to deliver the shipment ($T_{\text{transit}} = 16 \text{ hours}$). Finally, let's factor in 48 hours of overhead time ($T_{\text{overhead}} = 48 \text{ hours}$).

Plugging these values into Equation 1 and plotting the data-transfer time versus data size using the Ship It! strategy produces the maroon line in Figure 3. For the sake of comparison, the blue line shows the data-transfer time of the Upload It! strategy using a fiber-based Internet connection running at 800Mbps sustained upload rate. The figure shows

comparative growth in data-transfer time between uploading at 800Mbps versus copying it to LTO-6 tapes and shipping it overnight.

Equation 1 shows that a significant amount of time in the Ship It! strategy is spent copying data into and out of the storage appliance. The shipping time is comparatively small and constant (even if you are shipping internationally), while the drive-to-drive copy-in/copy-out time increases to a very large value as the size of the content being transferred grows. Given this fact, it's hard to beat a fiber-based connection on raw data-transfer speed, especially when the competing strategy involves copy in/copy out using an LTO-6 tape drive running at 160MBps.

Often, however, you may not be so lucky as to have access to a 1Gbps upload link. In most regions of the world, you may get no more than 100Mbps, if that much, and rarely so on a sustained basis. For example, at 100Mbps, Ship It! has a definite advantage for large data volumes, as in Figure 4, which shows comparative growth in data-transfer time between uploading at 100Mbps versus copying the data to LTO-6 tapes and shipping it overnight.

The maroon line in Figure 4 represents the transfer time of the Ship It! strategy using LTO-6 tapes, while this time the blue line represents the transfer time of the Upload It! strategy using a 100Mbps upload link. Shipping the data using LTO-6 tapes is a faster means of getting the data to the cloud than uploading it at 100Mbps for data volumes as low as four terabytes.

What if you have a much faster means of copying data in and out of the storage appliance? How would that compete with a fiber-based Internet link running at 800Mbps? With all other parameter values staying the same, and assuming a drive-to-drive copy-in/copy-out speed of 240MBps (50% faster than what LTO-6 can support), the inflection point (that is, the content size at which the Ship It! strategy becomes faster than the Upload It! strategy at 800Mbps) is around 132 terabytes. For an even faster drive-to-drive copy-in/copy-out speed of 320MBps, the inflection

point drops sharply to 59 terabytes. That means if the content size is 59TB or higher, it will be quicker just to ship the data to the cloud provider than to upload it using a fiber-based ISP running at 800Mbps.

Figure 5 shows the comparative growth in data-transfer time between uploading it at 800Mbps versus copying it at a 320MBps transfer rate and shipping it overnight.

This analysis brings up two key questions:

- If you know how much data you wish to upload, what is the minimum sustained upload speed your ISP must provide, below which you would be better off shipping the data via overnight courier?

- If your ISP has promised you a certain sustained upload speed, beyond what data size will shipping the data be a quicker way of hauling it up to the cloud than uploading it?

Equation 1 can help answer these questions by estimating how long it will take to ship your data to the datacenter. This quantity is $(\text{Transfer Time})_{\text{hours}}$. Now imagine uploading the same volume of data (V_{content} Megabytes), in parallel, over a network link. The question is, what is the minimum sustained upload speed needed to finish uploading everything to the datacenter in the same amount of time as shipping it there. Thus, you just have to express Equation 1's left-hand side $(\text{Transfer Time})_{\text{hours}}$ in terms of the volume of data (V_{content} Megabytes); and the required minimum Internet connection speed ($\text{Speed}_{\text{upload}}$ Mbps). In other words: $(\text{Transfer Time})_{\text{hours}} = 8 \times V_{\text{content}} / \text{Speed}_{\text{upload}}$.

Having made this substitution, let's continue with the scenario: LTO-6-based data transfer running at 160MBps, overnight shipping of 16 hours, and 48 hours of overhead time.

Figure 6. Minimum necessary upload speed for faster uploading.





Figure 7. Maximum possible data size for faster uploading.



Also assume there is 1TB of data to be transferred to the cloud.


The aforementioned substitution reveals that unless the ISP provides a sustained upload speed ($Speed_{upload}$) of at least 34.45Mbps, the data can be transferred faster using a Ship It! strategy that involves an LTO-6 tape-based data transfer running at 160MBps and a shipping and handling overhead of 64 hours.

Figure 6 shows the relationship between the volume of data to be transferred (in TB) and the minimum sustained ISP upload speed (in Mbps) needed to make uploading the data as fast as shipping it to the datacenter. For very large data sizes, the threshold ISP upload speed becomes less sensitive to the data size and more sensitive to the drive-to-drive copy-in/copy-out speeds with which it is competing.


Now let's attempt to answer the second question. This time, assume $Speed_{upload}$ (in Mbps) is the maximum sustained upload speed that the ISP can provide. What is the maximum data size beyond which it will be quicker to ship the data to the datacenter? Once again, recall that Equation 1 helps estimate the time required ($(Transfer\ Time)_{hours}$) to ship the data to the datacenter for a given data size ($V_{content}$ MB) and drive-to-drive copy-in/copy-out speeds. If you were instead to upload $V_{content}$ MB at $Speed_{upload}$ Mbps over a network link, you would need $8 \times V_{content} / Speed_{upload}$ hours. At a certain threshold value of $V_{content}$, these two transfer times (shipping versus upload) will become equal. Equation 1 can be rearranged to express this threshold data size as illustrated in Equation 2.

Figure 7 shows the relationship between this threshold data size and the available sustained upload speed from the ISP for different values of drive-to-drive copy-in/copy-out speeds.

Equation 2 also shows that, for a given value of drive-to-drive copy-in/copy-out speed, the upward trend in $V_{content}$ continues up to a point where $Speed_{upload} = 8 / \Delta T_{data\ copy}$, beyond which $V_{content}$ becomes infinite, meaning it is no longer possible to ship the data more quickly than simply uploading it to the cloud, no matter how gargantuan the data size. In this case, unless you switch to a



Data can be transferred at various levels of granularity such as logical objects, buckets, byte blobs, files, or simply a byte stream.



faster means of copying data in and out of the storage appliance, you are better off simply uploading it to the destination cloud.

Again, in the scenario of LTO-6 tape-based data transfer running at 160MBps transfer speed, overnight shipping of 16 hours, and 48 hours of overhead time, the upload speed beyond which it's always faster to upload than to ship your data is 640Mbps. If you have access to a faster means of drive-to-drive data copying—say, running at 320Bps—your ISP will need to offer a sustained upload speed of more than 1,280Mbps to make it speedier for you to upload the data than to copy and ship it.

Cloud-to-Cloud Data Transfer

Another strategy is to transfer data directly from the source cloud to the destination cloud. This is usually done using APIs from the source and destination cloud providers. Data can be transferred at various levels of granularity such as logical *objects*, *buckets*, *byte blobs*, *files*, or simply a *byte stream*. You can also schedule large data transfers as batch jobs that can run unattended and alert you on completion or failure. Consider cloud-to-cloud data transfer particularly when:

- ▶ Your data is already in one such cloud-storage provider and you wish to move it to another cloud-storage provider.
- ▶ Both the source and destination cloud-storage providers offer data egress and ingress APIs.
- ▶ You wish to take advantage of the data copying and scheduling infrastructure and services already offered by the cloud providers.

Note that cloud-to-cloud transfer is conceptually the same as uploading data to the cloud in that the data moves over an Internet connection. Hence, the same speed considerations apply to it as explained previously while comparing it with the strategy of shipping data to the datacenter. Also note that the Internet connection speed from the source to destination clouds may not be the same as the upload speed provided by the ISP.

Cost of Data Transfer

LTO-6 tapes, at 0.013 cents per GB,¹⁸

provide one of the lowest cost-to-storage ratios, compared with other options such as HDD or SSD storage. It's easy to see, however, the total cost of tape cartridges becomes prohibitive for storing terabyte and beyond content sizes. One option is to store data in a compressed form. LTO-6, for example, can store up to 6.25TB per tape¹⁸ in compressed format, thereby leading to fewer tape cartridges. Compressing the data at the source and uncompressing it at the destination, however, further reduces the effective copy-in/copy-out speed of LTO tapes, or for that matter with any other storage medium. As explained earlier, a low copy-in/copy-out speed can make shipping the data less attractive than uploading it over a fiber-based ISP link.

But what if the cloud-storage provider loaned the storage appliance to you? This way, the provider can potentially afford to use higher-end options such as high-end SSDs or a combination HDD-SSD array in the storage appliance, which would otherwise be prohibitively expensive to purchase just for the purpose of transferring data. In fact, that is exactly the approach that Amazon appears to have taken with its Amazon Web Services (AWS) Snowball.³ Amazon claims that up to 50TB of data can be copied from your data source into the Snowball storage appliance in less than one day. This performance characteristic translates into a sustained data-transfer rate of at least 600MBps. This kind of a data-transfer rate is possible only with very high-end SSD/HDD drive arrays with read-ahead caches operating over a fast interface such as SATA Revision 3, SAS-3, or PCI Express, and a Gigabit Ethernet link out of the storage appliance.

In fact, the performance characteristics of AWS Snowball closely resemble those of a high-performance NAS (network-attached storage) device, complete with a CPU, on-board RAM, built-in data encryption services, Gigabit Ethernet network interface, and a built-in control program—not to mention a ruggedized, tamper-proof construction. The utility of services such as Snowball comes from

the cloud provider making a very high-performance (and expensive) NAS-like device available to users to “rent” to copy-in/copy-out files to the provider's cloud. Other major cloud providers such as Google and Microsoft are not far behind in offering such capabilities. Microsoft requires you to ship SATA II/III internal HDDs for importing or exporting data into/from the Azure cloud and provides the software needed to prepare the drives for import or export.¹⁶ Google, on the other hand, appears to have outsourced the data-copy service to a third-party provider.⁸

One final point on the cost: unless your data is in a self-managed datacenter, usually the source cloud provider will charge you for data egress,^{4,5,12,15} whether you do a disk-based copying out of data or cloud-to-cloud data transfer. These charges are usually levied on a per-GB, per-TB, or per-request basis. There is usually no data ingress charge levied by the destination cloud provider.

Conclusion

If you wish to move big data from one location to another over the Internet, there are a few options available—namely, uploading it directly using an ISP-provided network connection, copying it into a storage appliance and shipping the appliance to the new storage provider, and, finally, cloud-to-cloud data transfer.

Which technique you choose depends on a number of factors: the size of data to be transferred, the sustained Internet connection speed between the source and destination servers, the sustained drive-to-drive copy-in/copy-out speeds supported by the storage appliance and the source and destination drives, the monetary cost of data transfer, and to a smaller extent, the shipment cost and transit time. Some of these factors result in the emergence of threshold upload speeds and threshold data sizes that fundamentally influence which strategy you would choose. Drive-to-drive copy-in/copy-out times have enormous influence on whether it is attractive to copy and ship data, as opposed to uploading it over the Internet, especially when competing with an optical fiber-based Internet link. ■

Related articles on queue.acm.org

How Will Astronomy Archives Survive the Data Tsunami?

G. Bruce Berriman and Steven L. Groom
<http://queue.acm.org/detail.cfm?id=2047483>

Condos and Clouds

Pat Helland
<http://queue.acm.org/detail.cfm?id=2398392>

Why Cloud Computing Will Never Be Free

Dave Durkee
<http://queue.acm.org/detail.cfm?id=1772130>

References

1. Apple. Thunderbolt, 2015; <http://www.apple.com/thunderbolt/>.
2. Amazon Web Services. Global infrastructure, 2015; <https://aws.amazon.com/about-aws/global-infrastructure/>.
3. Amazon. AWS Import/Export Snowball, 2015; <https://aws.amazon.com/importexport/>.
4. Amazon. Amazon S3 pricing. <https://aws.amazon.com/s3/pricing/>.
5. Google. Google cloud storage pricing; <https://cloud.google.com/storage/pricing#network-pricing>.
6. Google. Cloud storage transfer service, 2015; <https://cloud.google.com/storage/transfer/>.
7. Google. Google fiber expansion plans; <https://fiber.google.com/newcities/>.
8. Google. Offline media import/export, 2015; <https://cloud.google.com/storage/docs/offline-media-import-export>.
9. Herskowitz, N. Microsoft named a leader in Gartner's public cloud storage services for second consecutive year, 2015; <https://azure.microsoft.com/en-us/blog/microsoft-named-a-leader-in-gartners-public-cloud-storage-services-for-second-consecutive-year/>.
10. SCSI Trade Association. Serial Attached SCSI Technology Roadmap, (Oct. 14, 2015); <http://www.scsita.org/library/2015/10/serial-attached-scsi-technology-roadmap.html>
11. IEEE. 802.3: Ethernet standards; <http://standards.ieee.org/about/get/802/802.3.html>.
12. Microsoft. Microsoft Azure data transfers pricing details; <https://azure.microsoft.com/en-us/pricing/details/data-transfers/>.
13. Ookla. America's fastest ISPs and mobile networks, 2015; <http://www.speedtest.net/awards/us/kansas-city-mo>.
14. PCI-SIG. Press release: PCI Express 4.0 evolution to 16GT/s, twice the throughput of PCI Express 3.0 technology; http://kavi.pcisig.com/news_room/Press_Releases/November_29_2011_Press_Release_/.
15. Rackspace. Rackspace public cloud pay-as-you-go pricing, 2015; <http://www.rackspace.com/cloud/public-pricing>.
16. Shahan, R. Microsoft Corp. Use the Microsoft Azure import/export service to transfer data to blob storage, 2015; <https://azure.microsoft.com/en-in/documentation/articles/storage-import-export-service/>.
17. The Serial ATA International Organization. SATA naming guidelines, 2015; <https://www.sata-io.org/sata-naming-guidelines>.
18. Ultrium LTO. LTO-6 capacity data sheet; http://www.lto.org/wp-content/uploads/2014/06/ValueProp_Capacity.pdf.
19. Ultrium LTO. LTO-6 performance data sheet; http://www.lto.org/wp-content/uploads/2014/06/ValueProp_Performance.pdf.
20. USB Implementers Forum. SuperSpeed USB (USB 3.0) performance to double with new capabilities, 2013; <http://www.businesswire.com/news/home/20130106005027/en/SuperSpeed-USB-USB-3.0-Performance-Double-Capabilities>.

Sachin Date (<https://in.linkedin.com/in/sachindate>) looks after the Microsoft and cloud applications portfolio of e-Emphasys Technologies. Previously, he worked as a practice head for mobile technologies, an enterprise software architect, and a researcher in autonomous software agents.

Copyright held by author.
 Publication rights licensed to ACM \$15.00.