

Moving Data is Expensive

Two hundred years ago it would take months to transport a letter between or across continents. The further the letter had to travel, the longer it would take for it to get there. Then came the telephone, and now, instead of sending a letter, one could place a phone call and communicate the information in real-time. The connectedness of the world and the economies and the increasing dependency on information to deliver services, make decisions, and create value have placed heavy demands on data access. Achieving data access using traditional approaches, such as moving or copying data to where compute is, has associated costs, complexity, and consequences. Organizations attempt to address the symptom of the problem “how to move more data faster over distance,” but what they should be asking is “how to achieve real-time data access in support of applications without moving or copying data?”

The Cost of Moving Data

There are three main sources of costs when moving or copying data across distance. These are: cost of underutilized network, the cost of productivity loss due to time required to move or copy data, and the cost of storage used to store duplicates of data in multiple locations.



Underutilized Network

It is common to deploy a 1 Gigabit or 10 Gigabit network connection between two sites. The longer the distance between the two sites, the more latency can be expected. Typically, an application using storage located in the same datacenter may expect latency in the single milliseconds, but when reaching over distance, latency can run in the tens and hundreds of milliseconds. The greater the latency, the more throughput rates get affected.

Example: it is not unusual to see only ~6 MBps throughput on a Gigabit connection (maximum measured at 12 MBps). The result is that though one gigabit Ethernet connection is sized to move 450 GB per hour, the actual throughput ends up moving only 43.2 GB per hour. If the requirements is to move 10 TB across 3,000 miles, even with latency less than 100 milliseconds, it would take 9.6 days to move it. The other way to think about it is if a gigabit Ethernet connection costs you \$1,000 per month and you use less than 10% of bandwidth available, instead of the price per gigabyte moved being \$1,000/total capacity of gigabytes can be moved = \$0.00309 per gigabyte, the cost is \$1,000/ability to send gigabytes = \$0.032 per gigabyte.

	1 Gbps without Vcinity	1 Gbps with Vcinity
GB moved per hour:	43.2 GB	403.2 GB
Cost per GB moved: Assume monthly cost of 1 Gbps connection is \$1,000:	\$0.32 per GB moved \$32/TB	\$0.0034 per GB moved \$3.40/TB

Loss in Productivity



Data created outside of the datacenter must be made available to the applications in order to gain insights or create products and services. Every time data must be moved, the organization incurs a productivity penalty. This is the time that work is not being done. In some industries, Media and Entertainment being one, professional editors, colorists, special effects artists may sit idle waiting for data to become available. It is exacerbated further if the artists are in different cities, countries, or even continents. Data may have to be either

moved multiple times, each time incurring the productivity penalty, or copies would be created for each location, exposing them to security risks.

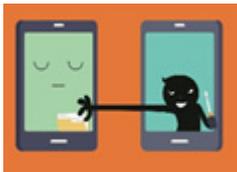
Example: Filming has completed, and the artists are awaiting video to begin work. It is a hot movie that needs to hit the theaters before the holidays. The studio must either deploy applications to do at least initial work in a mobile datacenter or move data to the application. The delay may be hours, but if the artists are being paid \$150 per hour and there are many of them, it adds up. The time delay also may affect when the final product will be ready.

Proliferation of Data Duplicates



Every time data is moved or copied, storage resources must be made available to store it. Whether it is persistent storage or a caching device, disk drives are deployed to catch data being sent over. If 10 TB are being sent, there is a need for at least 10TB of storage to be available at every site where this data is being sent. Standard market rates for storage are around \$300 per TB over a three-year time frame. If there is a single copy made, the 10 TB end up costing \$3,000 in storage. This is marginally accurate since it is challenging to procure small amounts of storage efficiently. The economy of scale, regarding storage capacity, kicks in at over 40 TB.

Liability and Security



If you are providing a service that requires the review and analysis of sensitive information, and the information must be moved to your facility, and you can send it over the wire because the liability of a data breach is in the millions, you wish you could process the data in place. But, since data must be moved, there are fees to insurance companies to protect against exposure, there is time that it takes to physically move data. This requires a storage system that is transported with the data. At the end of the process, the data must be destroyed. Conversely, if data that needs to be analyzed can't be moved at all due to regulations, compute resources must be moved to where the data is. In both cases, there is added complexity, liability, and cost associated with insuring, transporting, and analyzing sensitive data.

Summary

It is not enough to calculate the cost of acquisition, power and cooling, floor space, and administrative costs; determining the total cost of ownership requires the consideration of cost of opportunity. In financial institutions it is standard to calculate the opportunity cost of an investment, but in the enterprise, it is often too complex or cumbersome to determine what are the costs associated with an investment or lack of investment. It is clear that data mobility is expensive, not only in dollars spent per terabyte moved, but also the hidden costs such as cost of lost productivity, loss of revenue, diminished addressable market, and exposure to security risks. Though these are not glaring immediate costs, they may be deterministic long term. What are the consequences if one organization chooses to rely solely on mail versus another leveraging telephony?