Clipper Notes™
VENDOR-AGNOSTIC EXPLANATIONS AND ADVICE FOR THE INFORMATION TECHNOLOGY BUYER
Navigating Information Technology Horizons℠

Published Since 2007 Report #TCG2012026 November 8, 2012

Bringing "the Power" Closer to Home — Introspections for Future Storage Architectures

Analyst: Mike Kahn

The Problem

With a plug-in electric vehicle, there are no more trips to the gas station. Assuming that the whole idea of an electric vehicle makes sense for you (most importantly, that you are able to exist within the limits of its driving range), you get a great advantage of being able to "fill up" (charge your vehicle's batteries) at home. Sometimes, it just makes a lot of sense to do something closer to home, so that you don't have to wait or endure the transport time and inconvenience of going out for fuel or, more generally, dealing with whatever is your constrained resource. **Distance matters, primarily due to the time it takes to travel the distance.** Of course, speed of travel is another variable in this equation.

The same might be said for data, especially for "big data"[1]. **Putting data near the server cores that process and/or analyze it seems to makes a lot of sense (intuitively). Of course, there is the big problem of what to do when the data that you need is many (think "thousands") of miles or kilometers away.** Thus, there is an overwhelming tendency to transport the data to be closer to you and your local analytics servers (usually via FTP). This works when the data doesn't change much, when it isn't too big, and when this method isn't being done very often by too many people. Otherwise, you may have many or all of the problems that follow. **Again, distance matters, but often it is the cost of providing fast data access over that distance that is the root of the problem.**

First, in an era of running analytics on the most current (i.e., "real-time") data, getting a snapshot of a moving target always is a single-use happening (or should be). If you, or someone else, want to take another snapshot, the data likely will have changed (or been added to) and it will have to be "snapped" again. Each of these snaps has to be transported to you and stored (at least temporarily). If kept for longer than a few moments, you now have an out-of-date snapshot (or several or many) taking up space and containing potentially imperfect information (and often it is impossible to know what has changed and, thus, to know how imperfect it might be for your intended uses).

Second, if the data is big (think "really big", say measured in terabytes), just moving it to a distant location can be a heavy burden – on the network, on the receiving storage devices, and (of course) for the user who is waiting for the big data to be moved through what almost always is a constrained, shared pipeline. Yes,

IN THIS ISSUE

[1] There are many kinds of big data. Sometimes, it inherently is big (in terms of bits/bytes), such as high-resolution video (think of all of those security cameras streaming data 24 hours a day). Some of it is big because of the quantities of what is being stored (think about digital images in a large hospital). Some of it is big because of the large number of users creating and using it, think of Twitter or YouTube. However, in many cases, it is a combination of all of the above. If you don't know, inherently, what makes you the owner of the problems and challenges of big data (as measured in petabytes and larger), the phenomena and implications described herein may not apply to your world…yet. Remember that big data is not new. There just is so much more of it today, demand for analyses is growing rapidly, and "time-to-results" is the new metric.

networks are getting faster, but moving data long distance is never simple or free and constraints on available bandwidth often are a source of moderate-to-great delay (which might be measured in hours, days, or longer, as will be shown). Again, if speedy (think "instantaneous") availability is needed, the cost of the bandwidth to make this seem to happen instantly often is very high, especially for very large data sets.

Third, if you had to suffer the delay (and cost) of moving that big object through a too small (constrained) pipe, you may be inclined to operate on a "recent" copy, even as it becomes more stale with each passing minute (and hour and day). Not only are you now duplicating a snapshot or even making "just a copy", you begin to recognize that, especially with really big data, everything seems to be working against you. Here are some reasons why.

- **It takes too long to move the data to where you can use it.** While this might seem too strong of a conclusion, just multiply what you want to do by thousands of times per year (several times a day) or even millions of times (several thousand times per day), and you begin to see why most enterprise networks are not cut out to move massive data many thousands of times per day (or even per week, etc.), especially over a long distance. Whatever network bandwidth you have at your disposal, it probably isn't going to cut it as your data reaches petascale and exascale proportions, at least to your users' satisfaction. While we now have come to expect one gigabit per second LANs as the norm within a building or 10 gigabit per second LANs as the norm within a data center, you need to recognize that your slice of a long distance network (a.k.a. wide area networks (WANs) that move data between metro areas) most likely is moving at megabit speeds, at best. If you have to move a gigabyte file at a megabit per second, it would take a thousand seconds (about 2.8 hours). If you have to move a terabyte file, it would take about 16.5 weeks (around the clock). If you have to copy the contents of a 3 TB hard disk, it will take almost a year at this speed. As a point of reference, your Internet connection at home might be going at 10 megabits per second (but probably is far slower on larger files, especially on uploads). As a more generous example, if you wanted to move 50 gigabytes at 100 megabits per second (which is pretty fast), that would take 1.4 hours. (See Exhibit 1, at the top of the next page, for more on the math, the networking assumptions, and the equivalent transport times.)
- **It costs too much to move this kind of big data frequently.** If you need to analyze locally the remote 50 GBs just mentioned, you might have to continually pre-stage the movement of data across the country, say every 90 minutes. If you only have 50GBs of data, maybe you could live with this. However, if this is just one of many data sets being used (say, of maybe 100 TBs total or more), you might have more to transfer regularly than you can economically justify, especially for data that always will be about 90 minutes old.[2] If you have petabytes of data that you have accumulated, just moving/using a little of it "on demand" will have large implications, in terms of time and cost.
- **It takes a lot of space (and money) to store the now duplicated data.** In addition, you now likely are paying to back it up at multiple locations, especially if there are many others like you who are asking for similar extracts of data or want to see the same large files as you do. Redundant data is a heavy cost multiplier.[3]
- **Even worse, at the end of the day, the data users begin to realize how much of their own time they are spending** just to get the data to your local computers to do the analysis that is (only) an accelerant to their primary objective, which is to make a decision or recommendation or to draw some insightful conclusions and make further progress toward something that is important.

Fourth, as said above, distance matters. If you have to move big data across the country, it is easy to imagine why big distances matter, in terms of time and cost. However, little distances matter as well, as will be explained. In this case, short distances are measured in meters or feet (or maybe a couple of kilometers or miles) but maybe much less. For a vivid example, think of the distances within a data center (although it equally might be true within a metro area). Two shorter distances are worth considering, but the concept that "distance matters" applies to them both (as well as the long distances discussed above).

---

[2] Because it takes that long (1.4 hours) to complete the transfer.

[3] However, if you are using streamed data on a "single use" basis, this may not apply.

**Exhibit 1 —
How Long It Will Take To Transport Different Volumes of Data?**

Data Size		at 1 megabit/second		at 10 megabits/second		at 50 megabits/second		at 100 megabits/second		at 1 gigabit/second	
1	GB	2.8	hours	0.3	hours	0.1	hours	1.7	minutes	0.2	minutes
10	GB	1.2	days	2.8	hours	0.6	hours	16.7	minutes	1.7	minutes
50	GB	5.8	days	13.9	hours	2.8	hours	1.4	hours	8.3	minutes
1	TB	16.5	weeks	1.7	weeks	2.3	days	1.2	days	2.8	hours
3	TB	49.6	weeks	5.0	weeks	6.9	days	3.5	days	8.3	hours
100	TB	31.7	years	3.2	years	33.1	weeks	16.5	weeks	1.7	weeks

Notes
(a) All numbers above are rounded to the significant digit displayed.
(b) In the above calculations, it is assumed that there are 10 bits in a byte (rather than 8). This is done to accommodate for parity checking and other networking overhead. (In addition, it is mathematically convenient.) In comparison to the factors not considered (see below), this effect is negligible.
(c) In the real world, you may only get 1/3 to 2/3 of the rated speed of a WAN, due to traffic, collisions, overhead, etc. This has NOT been factored in.
(d) In the above durations, distance is not considered. Of course, it takes more technology and budget to deliver the transfer rates over a long distance. While gigabit speeds theoretically are possible over a short distance, say between your laptop and a nearby server, getting data from a distant cloud likely will take place at speeds under 10 megabits per second (and often at less than 1 megabit per second), especially if the files are very large.
(e) As a point of reference, a gigabit/second of bandwidth is equal to a modest optical Internet backbone connection.
(f) As a point of reference, 3TB is the capacity of a large hard disk.

Source: The Clipper Group, Inc.

- **The distances between data center devices.** This might be described as a couple to several meters. Think of the distance between a server and a storage array.
- **The distances within a device.** This might be measured in centimeters or inches, but may be much less (e.g., the small distances between the memory and CPU within a server).

If distance matters regardless of whether it is across the country or across a motherboard, does it matter the same? Also, is distance the right metric for measuring the affect? Let's take the second question first. No, linear distance is not the best unit of measure. Physical distances helps us imagine why things are the way they are (in the same way that Newtonian physics help us to comprehend the world around us) but when the distances get very small (for example, within a processor chip), the point of relativity (to our physical world) drops beyond our usual vehicles for comprehension (as also happens with quantum physics). **What really matters (and is imaginable even on a very small scale) are time and money.**

We easily can comprehend time that is measured in hours, minutes and seconds. We can even comprehend small fractions of a second, usually by reference to real-world examples. We know, intuitively, when something is slow (like when we Google something and it seems to take forever, for whatever reason). Either it comes back quickly (i.e., within our expected perception of "very fast") or seems to take "too long". If we are downloading a 200MB update from Microsoft for *Windows* or *Office*, we intuitively know whether that is going quickly or seems to be taking forever.

Thus, all of this is relative to our expectations. While most of us can't fathom the difference of going from 2.0 to 2.5 GHz per second CPU speed (since we don't have a good reference point for billionths of a second), we still can fathom (at a gross level) what a 25% improvement in CPU speed might mean for us as application users. **Unfortunately, when it comes to analytics (especially), expectations are rising.**

The bottom line is that time to arrival (or completion) really matters, whether you are waiting for a package from Amazon to come from afar, hot water to reach your faucet from a distant hot water heater, or data to move to the core that will process it. While *time to arrival* is an important metric, especially if you, personally, are the one doing the waiting, the *cost of making this happen* likely also is an important concern, if not to you then to the one paying for the delivery.

A Eureka! Moment

You already might have had your "Eureka! Moment" (where that light bulb in your head goes on). If not, let me help you.

Why, you might be asking yourself, do I want to move large data sets across the country or around the world, especially data that is constantly changing and whose value decreases quickly as it ages?[4]

Let's take a couple of steps backwards and consider this.

First, you may be motivated by a need to do some kind of business or statistical analysis (these days, this activity often is called "analytics"). You may be looking for patterns, insight, recommendations or, even, a decision. You have a process or methodology, whether canned in advance or just created for a single use, which you want to apply using your collection of analytical tools. This may be an iterative or cumulative process requiring many passes through the data, with you reviewing the interim results and drilling down to get the answers that you seek. Because you know that this kind of analysis can be both processor and I/O intensive[5], you usually strive to do it locally (as in "closer to you, i.e., near *your* local servers). This has been the prevailing thinking for many decades, but is seriously flawed for really "big data" in a world of distributed users and multiple data centers, including data sitting somewhere in a cloud on the Internet.

Hopefully, your light bulb is glowing. Locality is critical but unless you have just one data center and all of your "analysts" are sitting "nearby" (in the same metro area), there is a good chance that that the data you need won't be nearby (at least not without all of the copying and transferring discussed above) and, thus, not quickly available without spending dearly to make enough bandwidth available whenever you need it. **Moving data and processing it locally makes sense only if the data is nearby (or if the data does not change and you will be analyzing it for more than a short time).** Let me restate the key thought here: ***Processing it locally makes sense only if the data is nearby!***[6]

If moving the big data is akin to shoving a brick down a garden hose, why are you so set on moving it? **Why don't you operate on it (i.e., analyze it) where it sits and just return the answers, i.e., patterns, insight, recommendations or, decision?** The light bulb now should be glowing brightly.

None of This Thinking is New

When I was an industrial engineering student many decades ago, this was a common transportation/distribution/warehousing metaphor, loosely called "The Transportation Problem"[7]. In the digital age, we have seen many solutions for spreading data to many locations across the Internet to speed up delivery of content (because it is nearby or closer). So this Eureka! Moment is not new for information systems either.

What is new is the bigness of the problem or, more specifically, the *bigness of the data* and the fact that *speedy bandwidth is rarely without a high cost*. In an era of terabyte-sized data collections (or larger), the volume or the amount of data being moved (i.e., "transported") now is exceedingly "heavy" in terms of what it takes to "lift" it and get it to where it is needed. Moving something heavy at high speed always is costly.[8] **Data transport is made even more challenging because the *currency* of the data is what drives the most successful uses of analytics.** Just because you lifted it yesterday, doesn't mean that you don't have to do it several (or many) more times today, just to include whatever has changed or been added or to meet some new request(s).

The Eureka! conclusion is that *data should be analyzed near where it sits*. However, different images come to mind on what this means and how it might be implemented. A typical first image might have data in an array in the same data center as the server core(s) that are operating on it, with connectivity provided by Fibre Channel or Ethernet[9]. The array might have solid-state disks as a primary tier or as a caching accelerator. Regardless of the speed of the storage, the entire volume of data still has to be transported to the

---

[4] As has been suggested, the issues can the same for much shorter distances. More on this, to follow.

[5] Computational analytics (e.g., mathematically-intensive modeling and correlation) tends to be processor intensive when there is little waiting to access the data. With processing power being plentiful these days, it is waiting for data that tends to be the bottleneck.

[6] "Nearby" is a relative term, as you will see.

[7] The Transportation Problem classically dealt with determining how many warehouses you needed to house your inventory and where they should be located. The goal was to minimize warehousing costs, inventory levels, handling, and transportation costs while maximizing customer satisfaction (usually measured in terms of time to delivery). This kind of multivariable optimization was difficult to solve (too many variables for simultaneous optimization) and produced continuing headaches. Not much has changed.

[8] Just think about what it takes (in terms of energy and costs) to put large object, like a satellite or space station module, into orbit.

[9] Or FICON, for mainframes.

server over the likely shared Fibre Channel or Ethernet. Regardless of whether the Fibre Channel is running at 8 or 16 gigabits per second[10] or Ethernet at 1 or 10 gigabits per second[11], at some point the weight of what is being transported (often called "the transport volume") becomes a burden and then becomes a bottleneck. There many reasons why, but two are worth noting.

1. **The communications network and array(s) are being shared and many applications and users are making simultaneous demands.** As the number of users and applications rise and the size of the data grows, the volume to be moved will become burdensome at some point, and usually more costly to deliver speedily, as more expensive technologies need to be deployed more widely to cope with the ever-increasing demands. There may be many switches or directors through which the data must pass as it moves down the network.[12] This is why maximum rated ("peak") network speeds rarely are attained (except for an occasional burst).
2. **The server has to wait for the data to arrive.** No matter how many cores you have to do your analysis, little can be done without accessing the data. Not only does the data have to be transported to the server (whatever the distance), it also has to be transported to the CPU core. To do that, it has to go through another network controller (sitting on or in the server) and through the server's bus. While this happens in small fractions of a second, it still is going at a speed that usually causes the central processor to have to wait.[13]

As seems to be happening regularly, we have gone full circle. In the old days (of decades ago), keeping the server busy was the highest priority (because it was so costly and the costs had to be recovered) and the users were told to "be patient", as they waited for results. **Now, data use (i.e., leveraging data through analytics) reigns supreme and data centers are scrambling to make this happen.**[14] **Nonetheless, making server cores wait means that the real bottleneck is downstream, *near where the data is stored*. If we can make the server wait less (by speeding access to needed data, maybe by putting the processing power closer to the data), the answers will come quicker, which is the name of the game at most enterprises.**

Implications for Servers and Storage

Here's a summary of key thoughts, so far.

- **Moving lots of data to where it is to be analyzed usually is burdensome and expensive over long distances.**
- **While everything happens faster within the same physical data center, data networks and storage arrays still cause the servers to wait (especially when the data volume is large or being read many (think "thousands" or more) times per day).**
- **Even short distances seem to matter (as measured in "time to results").**

Having all of your data in or very near the server's memory will speed the analytics processing. If I/O doesn't need to cross the motherboard's bus, everything will go faster than if it does and most certainly this will be faster than going out to an outboard array, which will be many times faster than moving the data from a distant location for local processing. **Thus, having data really close to the analytics' server core(s) might really speed the delivery of the answers that are being sought.** Whether by clever hardware and firmware design (where "memory" might be one of several levels of local cache) or other tricks of the trade, **keeping the data nearby can remove or reduce important bottlenecks that come from big data that is being stored "farther away".**

However, the flip side also may be worth considering. If the data is someplace else (far away), it may be better to send your questions to server cores near where the data lives and just have your answers returned to you. Yes, this sounds like time-sharing of eras gone by (if you are old enough to remember), but now is to

[10] Today's fastest Fibre Channel moves at a maximum speed of 16 gigabits per second.

[11] 10 GigE increasingly is the standard for Ethernet within the data center.

[12] You also need to consider the delays of going through networking cards or ports and seek times.

[13] Even if a solid state cache or disk is placed on the server's bus, it still has to endure the movement of the data across the network, through the network controller and through the bus). Only when the data has been "adequately cached" to such a solid state device, might the server be able to begin its analysis.

[14] This is true whether you are talking about an in-house data center or one sitting "in the cloud". The problem is the same.

considered to be the heart of a cloud strategy. If your data and processing is far away, you probably should analyze it remotely (from your own perspective), in order to shorten the time to get to your analytical objective.

Conclusion

There are many ways that this acceleration might be achieved. Exploring this in technical detail will be deferred to a later Clipper paper (or several), as this paper has been focused on *turning on your light bulb* to the nature of this growing "big data analytics" problem. There are three conclusions that you should take away.

1. **Being able to ask questions (i.e., run analyses) on large data sets likely will be done more quickly and with less expense if you transport your queries or structured analyses to where the data resides and only the answers are returned.** This is especially true when operating on frequently-changing "live data".
2. **Getting your data closer to the server cores (think "inside the server") will accelerate analytics and business intelligence even more.**
3. **Cost is a driving factor in the quest for speedy answers. If you want to ask a lot of questions regularly, you must consider what drives up your costs and how to keep them down.**

The purpose of this paper was to turn on the light to see that regularly moving very large amounts of data (your "transportation problem") will, inevitably, become a very serious operational bottleneck and/or cost overrun. This especially is true when long distances are involved but also true for the seemingly short distances within the data center. Now that you have reached this conclusion, you are ready to consider your architectural alternatives in a better light, which Clipper will do in future bulletins.

Today's Eureka! Moment is for you, whether you see the world as data center infrastructure or as an operational analytics challenge. Think about it! Ask yourself the tough questions about why you do what you do before you run out of time and/or money.

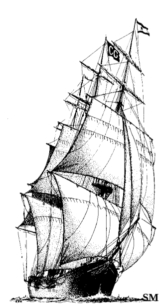

About The Clipper Group, Inc.

The Clipper Group, Inc., now in its twentieth year, is an independent publishing and consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

***Mike Kahn* is Managing Director and a cofounder of The Clipper Group.** Mr. Kahn is a veteran of the computer industry, having spent more than four decades working on information technology, spending the last 17 years at Clipper. For the vendor community, Mr. Kahn specializes on strategic marketing issues, especially for new and costly technologies and services, competitive analysis, and sales support. For the end-user community, he focuses on mission-critical information management decisions. Prior positions held by Mr. Kahn include: at International Data Corporation - Director of the Competitive Resource Center, Director of Consulting for the Software Research Group, and Director of the Systems Integration Program; at Power Factor Corporation, a Boston-based electronics startup – President; at Honeywell Bull - Director of International Marketing and Support; at Honeywell Information Systems - Director of Marketing and Director of Strategy, Technology and Research; at Arthur D. Little, Inc. - a consultant specializing in database management systems and information resource management; and, at Intel Corporation, Mr. Kahn served in a variety of field and home office marketing management positions. Earlier, he founded and managed PRISM Associates of Ann Arbor, Michigan, a systems consulting firm specializing in data management products and applications. Mr. Kahn also managed a relational DBMS development group at The University of Michigan, where he earned B.S.E. and M.S.E. degrees in industrial engineering.

- ***Reach Mike Kahn via e-mail at Mike.Kahn@clipper.com or via phone at (781) 235-0085 Ext. 121. (Please dial "121" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, **The Clipper Group Explorer**, **The Clipper Group Observer**, **The Clipper Group** *Captain's Log*, **The Clipper Group Voyager**, **Clipper Notes**, and *"clipper.com"* are trademarks of The Clipper Group, Inc., and the clipper ship drawings, *"Navigating Information Technology Horizons"*, and *"teraproductivity"* are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosures

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.