THE CLIPPER GROUP
Captain's Log™

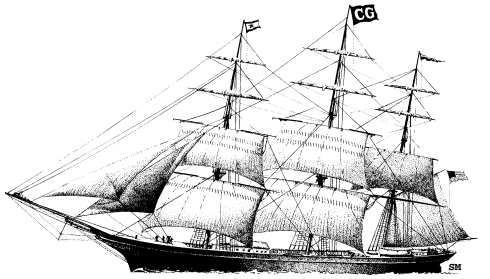

Navigating Information Technology Horizons℠

Published Since 2001 Report #TCG2007091 October 8, 2007

Roddenberry, Einstein and the Dinosaur — Considering the Unfathomable in IT Optimization

Analyst: Mike Kahn

The Star Trekian Challenge

It's the mid 1960s. Gene Roddenberry, creator of *Star Trek*, was creating a pilot for what became the fabled TV series and movies. His production budget was limited. How was he going to get the fictional crew to each of the worlds that they would visit? He couldn't afford the time it would take to create the miniature sets for the *Starship Enterprise* to make a landing on each weekly destination nor could he afford the cost of production. He needed a plausible, fast, and visually-credible transportation mechanism, so he "beamed" the crew to their encounter.[1] Now that's making the problem go away. If real life could only be that simple … sometimes, it is.

What CIO wouldn't want to have a "transporter" that would "beam" his infrastructure to where it needed to be? Sometimes, what seems like science fiction is possible. For example, you used to have to be a cutting-edge theoretical physicist to comprehend Einstein's *Theory of Special Relativity*, primarily because it was, to most, unfathomable.[2] What is known today about the universe was unfathomable to almost everyone two generations ago, unless you were a forwarding-looking theoretical astrophysicist. Yet today, it is explained in high-school science texts. Our point of reference has changed, and is contaminated by our preconceptions, as can be said about many things generational.

Science fiction has surrounded "the Mainframe" since the 1950s. Just consider what Spencer Tracy was promising in the movie *Desk Set* (1957) – answers to every question at your fingertips. Sixty years later, we call it Google. What was imaginary and unfathomable, even threatening to the corporate reference librarians headed by the character played by Katherine Hepburn, is accepted today as quite believable, if not quite 100% fulfilled. In this movie, you had to believe that the computer with the blinking lights, crunching noises, and spinning tapes that spit out punched cards with the right answer was fathomable. Many in the audience saw it as a prophecy of the future.

Today's Mainframe, IBM's *System z*9, has its roots in the IBM System/360, introduced in 1964.[3] Everybody thought that it could solve all the problems of the world. It looked not unlike the 1950s prop in Desk Set, minus the crunching sound effects, of course. IBM touted it as revolutionary and the "360" in its name connoted that it was applicable to many things (at all points on the compass). Few had the capacity to separate reality from the hype, because computers were unfathomable to all but a few. In reality, the S/360 began what was the business equivalent of "beaming" enterprises into a

---

[1] As told by Gene Roddenberry while on tour to an audience at the University of Michigan in the early 1970s, at which the author was present.

[2] The idea that the observer is some how involved in the problem (what is being observed) and that time and space are not absolute was (and still is) hard to comprehend. However, the analogy to the IT decision maker being involved with the "state of IT" and thus influencing his or her observations (or lack of observations) gets right to the heart of the issues discussed in this paper.

[3] See *The Beginning of I.T. Civilization - IBM's System/360 Mainframe* in The Clipper Group Captain's Log dated March 30, 2004, and available at http://www.clipper.com/research/TCG2004028.pdf.

fantastical future state, although we are still trying to live up to many of the promises made by IT organizations along the way. In the last four decades, we have accepted as fathomable what computers systems can do. That is good, because what is unfathomable is hard to project into the future, whether as technology or benefits to the enterprise.

Unfortunately, to many if not most folks during the 1960s through the 1980s, the Mainframe was unfathomable. Unless you studied computer science or electrical engineering, you might have a lot of trouble understanding *how it did what it did*. You might have understood *what it did*, but not how. There was no consumable point of reference. Even with minicomputers in the late 1970s and 1980s, most could not get their arms around a smaller version of "a big box." We just had to accept that the data, somehow, was beamed through some kind of programming inside the big box. The good news, for many, is that things seemed to work.

Then, along came the first PCs (Apple II and IBM PC). Suddenly, many more could see what a computer was, both inside and out. Over the last 25 years, most of us see the PC as the fathomable representation of *what a computer really is*; more specifically, *what they think a computer really is*.

This has led to a host of problems. In the 1980s and 1990s, "big computing" was said to be dead, there were no problems that could not be solved if one networked together enough PCs. IBM's Mainframe was declared a *dinosaur*, even by executives from IBM. Everything it did was considered *legacy*, as in *legacy applications* and *legacy thinking*. The Mainframe, workhorse of the world's largest enterprises, while fashionable for several decades, was never fathomable. That was - and still is - a big problem, not just for IBM but also for many enterprises who chose to ignore its potential to accelerate the visions of the future. Additionally, today's "industry standard server" is approaching, architecturally, the complexity of decades-old Mainframes.

If there were a trite and now-dated catch phrase to describe the situation for today's Mainframe, it might be *Not your father's Oldsmobile*.[4] The youngest readers will ask, *what's an Oldsmobile?* Others will recall it as another "big-box dinosaur" from a bygone era that died an uncomely death. Either there is no point of reference or the point of reference is too simplistic. Some might say that this is bad analogy, but today's mainframe is not yesterday's dinosaur. If your only point of reference is some archaic snippet, it is time to look a lot closer.

To appreciate what a Mainframe is, and what other computers systems still have yet to become, is a challenge, one that requires a certain amount of "mental beaming." You need to get to a place that you can't quite comprehend fully, if at all, and you've got to accept as reasonable the possibility that the Mainframe can actually get you there.[5] Once you accept that you can get there (there is a lot of substantial evidence), you will be in a much better position to ask, if you are so inclined, *how is this possible?*

Today's IBM Mainframe is not the mainframe of the 1980s or 1990s. IBM is not that heavy-leaning, proprietary systems vendor that some of us remember all to well. Unfortunately, Mainframe recollections and misconceptions have evolved into a whopper-sized mythology, which is seen as the truth, by many. After decades of being out of favor, it is accepted by many that largely centralized computing is the answer to many of the challenges facing IT organizations today. No platform does this better that the Mainframe.

How can that be, I can hear the *Grasshopper*[6] ask? It is *so big*, so costly, and *so … unfathomable*? Its architecture is almost 50 years old. *How can it be what we all seek?* There is a short and simple answer, which will follow shortly in a few paragraphs. There is a very long and more detailed answer, which someday I may write, maybe in several parts.[7]

---

[4] Yes, I seem to be on a trip through old TV series, but for this one, there is the connection in that Leonard Nimoy, who played (Mr. Spock), was the voice in the commercial that featured this phrase.

[5] That's quite a lot to ask, I know. It took me a long time to first accept and then understand the Theory of Special Relativity, at an "amateur" level. Once I got there, all kinds of other previously unfathomable thoughts and realities are now plausible and even somewhat fathomable. Fortunately, I cut my teeth on an IBM 360/67, the first of IBM's "virtual machines". The "insides" were real to me…I didn't have to beam myself to some sci-fi imaginary place…I was there. (See reference in footnote #2.)

[6] One last TV analogy. In the series *Kung Fu*, the lead character was called "Grasshopper" by his Shaolin monk masters when he was a boy, because had so much to learn.

[7] In the meantime, see *Mainframe Mythologies Live On - Setting the Record Straight* in The Clipper Group Captain's Log dated May 23, 2006, and available at http://clipper.com/research/TCG2006038.pdf.

Einstein and Relativity

Albert Einstein was accepted for being brilliant, even though few educated contemporaries could understand the thinking surrounding his profundity. The Mainframe may be like Einstein, brilliant, but not fathomable; ultimately accepted for paving the way to the future, even though it was and is a future than few can really comprehend and explain. To most, to contemplate the power of the Mainframe is akin to contemplating the power of the universe. It may be more fathomable and more acceptable if you can, for a little while (maybe once a week), assume the scope of that problem away, as did Gene Roddenberry, and build your future upon on a reality-shattering premise. Beam yourself to that reality to learn what it means, for you and your enterprise. Then you might be prepared to ask, *how is this all possible?* Why did it take so long for me to see the light? Once there, you might choose to utter the words of the fictional captain of the *Enterprise* in *Star Trek – The Next Generation*, Jean-Luc Picard: ***Make it so!***[8]

The Short Explanation

All fairy tales and science fiction are based upon the reader's sense of reality. If it is too unbelievable, i.e., not grounded in something experiential or accepted as truth, it becomes unpalatable and unsuccessful as narrative.[9] Theoretical physics, including the Theory of Special Relativity, is hard to fathom because it is not based on the common experiences of Newtonian physics, i.e., the laws of gravity and motion that surround our existence. To explain the Mainframe requires a "Newtonian" or experiential explanation, which, for many or most these days, admittedly lack. It is through one's sense of reality that we can understand what we cannot easily experience or even fathom.

Let's jump to a fathomable future, one that might even be today. As users of well accepted or even expected (i.e., entitled) technology, most of us really have no idea how things work.[10] When we "Google" something, we have no comprehension of how the results presented to us are made to happen. When we send a text message or email, or even just charge something on a credit card, we really can't explain how it really works, nor do we care. We know that it works, thus the process is fathomable. Because it is works, we might contemplate and even be able to explain (correctly or incorrectly) how it might really work but, in the end, we just accept that it works.[11]

Today, we, at least those in computer science and IT communities, comprehend a meaning to the phrase *grid computing*. We envision hundreds or thousands of PCs or servers (i.e., "industrial-strength PCs") somehow networked together to provide the horsepower to solve the problems of the day. This modern view of grid computing is based on our tangible reference point…our PC.[12] We can imagine that the resources can be shared and managed to meet the needs of the enterprise. We can't imagine quite how this is done, but we know that it is possible. What we don't consider is whether this is the most efficient way to get the jobs done, but since hype is often equated with reality, it must be so.

Well, that might have been true … nah, it probably never was true … for most non-computational applications. Furthermore, today, with increasing high-core and thread count processors requiring that the "PCs" virtualize to increase their utilization and spread their higher costs, it is a whole new ballgame. What we have done, and rightly so from a *scale-out* perspective, is move a lot more complexity into the data center architecture, which is exactly what has been criticized (internally) in scale-up architectures, like the Mainframe. This has happened at the same time that the *scale-up* architectures have reduced their costs dramatically, when compared to a decade earlier.[13] (Historically

---

[8] If you can handle the pun: *Make IT So!*

[9] Fantasy, too, is based on human iconography. Hieroglyphics may be indiscernible without a point of reference or perspective.

[10] This is a century-old theme in science fiction. It continues to replay itself, because it contains an abundance of truth and concern.

[11] Just like my car, I can only imagine how it really works, even though I am a graduate engineer. I know that it works, because it gets me there. Unless it doesn't work, I rarely think about how it works.

[12] For a longer discussion on grids, see *All Nodes Are Not Equal - Thinking Differently About Grid (as We Used to Know It)* in the issue of Clipper Notes updated on May 17, 2007, and available at http://www.clipper.com/research/TCG2007064.pdf.

[13] What many do not understand is that vendors cannot make a long-term, cash-cow business out of selling computer systems that are not feature and price competitive. IBM has seen significant growth in new uses of its Mainframe, primarily by offering what enterprises need (a place to run open applications, like Linux and Java apps, including SAP and Oracle financial/ERP applications and IBM DB2, Oracle and other DBMSs) and at a price that is competitive with so-called "open systems" (Intel and AMD-based) solutions.

high costs are the other complaint often leveled at scale-up systems.)

What needs to be done is some honest "level setting" followed by some technical and TCO comparisons. This is more than can be considered in this paper, but here are a few salient guidelines to consider.

1. The system with the most virtualization probably is the superior solution architecturally.
2. The system with the best virtualization probably is the most efficient and, concomitantly, the most effective.
3. The system with the most flexibility in dynamically provisioning assets (think "virtual resources for virtual servers") probably is the superior solution operationally.
4. The system with the most extensible capacities, in its largest instance, probably can offer the lowest cost per unit of resources.
5. The system with the highest-availability may be required, even if it costs more to implement.
6. The system with the most connectivity within the system (as opposed to externally between system components) probably has the lowest networking costs.
7. The system with the most policy-driven automation probably has the lowest cost of administration and operation.
8. The system with the most "open" approach to applications may have the widest applicability in the larger enterprise.
9. The system with the best end-to-end security solution (for, say, PCI DSS) trumps all others.
10. The system with the fewest components is a lot easier. Why take on responsibility for planning, building, wiring, and connecting hundreds or thousands of servers together when you might be able to get a preassembled and test solution one a single box or a small number of boxes?
11. The system designed to use the least amount of energy while doing the most work (energy efficiency) probably is the right system to keep Data Center costs from escalating as the cost of energy rise year after year.

Notice what is not on the list. It doesn't say:

1. The system with the most common (i.e., fathomable) components is, necessarily, the best solution.
2. The systems with the most network cabling are inherently better.
3. The solution with the least expensive components is, necessarily, the one with the lowest TCO.
4. The system that does *Windows* is the best solution for all other environments.

I'm sure that you know where I am going with this. **IBM's Mainframe is the best in all of the most important places.** Yes, this is a very strong statement and requires extensive proof, much more that this bulletin will allow. **For now, you need to accept that this premise is credible.** Like The Theory of Special Relativity, you can't begin by understanding it all. It really helps if you can open your mind to believing that its revelations are possible. If you can accept the premise that it might be true, then you are ready to look into the details.

Conclusion

The bottom line to this fractured fairy tale is this: **if you cannot imagine that the Mainframe might be the best solution, it will never be your solution.** *Boldly go* where many, possibly including yourself, have been unwilling to go – ask whether a Mainframe might be the right vehicle to transport you to a better future.[14]

[14] And, check back for another episode in this continuing saga.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent 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 founder of The Clipper Group.** Mr. Kahn is a veteran of the computer industry, having spent more than three decades working on information technology, spending the last decade and a half 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; President of Power Factor Corporation, a Boston-based electronics firm; at Honeywell Bull - Director of International Marketing and Support; at Honeywell Information Systems - Director of Marketing and Director of Strategy, Technology and Research; with Arthur D. Little, Inc. - a consultant specializing in database management systems and information resource management; and, for 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 MikeKahn@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.

Disclosure

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.

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.