



Mainframe Mythology Lives On — Setting the Large Systems Agenda

Analyst: Mike Kahn

One More Time – Large Systems Thinking

When it comes to large systems, in general, and mainframes, in particular, this much is clear – there is a lot of misinformation out there, driven, on one hand, by a lack of familiarity or knowledge and, on the other hand, by a collection of enduring myths. The myths about mainframes seem to have much more to do with their “stickiness” (in terms of the myths’ ability to stick to whomever hears them) than the truth. Having been on this soapbox many times, it seems that there are two groups of folks out there, regarding large systems thinking. Simply put, either *you get it* or *you don't!* If you get it, you might find this *Captain's Log* amusing. If you don't get it, this issue is for you.

In either case, this paper is about getting to *20-20 vision* about meeting enterprise information processing needs. To do this, I will present 20 characteristics that describe what enterprises really require in IT infrastructure, without being architecture specific. This will allow us to communicate on a common ground. Then, I will present 20 myths about mainframes that need to be recognized as such. With both, you should have *20-20 vision* to focus clearly about your enterprise's IT workloads (processing requirements), and how you are going to get them done within policy guidelines. The mainframe isn't the only answer, but it is an answer worthy of your serious, unbiased consideration.

20 Critical Characteristics of Enterprise Systems

I seem to spend a lot of time on this subject these days. Maybe, it's because I have been known by many - for a long time - as a *mainframe advocate*.¹ OK, I admit to seeing stars when I think about the mainframe. However, I'd like to broaden this perception, a little, by restating my admiration for large-scale computer systems, which include the mainframe, but also include most “Scale-Up” systems, i.e., those that employ a sharable, SMP architecture². To do this, let me share with you the key descriptors or characteristics that I think define most large “enterprise systems”.

1. **Shared-asset IT infrastructure with a centralized point of control.** This means that IT assets are seen as a pool of resources that can be shared among various users, organizations, etc. Sharing permits *resource optimization*. Nothing in this statement requires that some or all of the resources be centralized physically, but that is common.

¹ See *The Clipper Group Captain's Log* dated March 30, 2004, entitled *The Beginning of I.T. Civilization – IBM's System/360 Mainframe*, available at <http://www.clipper.com/research/TCG2004028.pdf>.

² *Symmetric Multi-Processor* architectures also can be found in large-scale *Unix* and *Linux* systems and in some *Windows* systems. Traditionally, they contain many processors (8 to 64 or more), typically within one or a few enclosures. The key differentiator is that the operating software subdivides the system into flexible partitions that allow system resources (processors, memory, storage, bandwidth, etc.) to be shared among applications (and users) that are running simultaneously. This is different from “Scale Out” architectures that tend to have rack or blade units that contain up to four processors, each within its own enclosure. Scaling is done by adding rack or blade units, hence the name “Scale Out”. Processing resources are not shared within these building blocks, although they may be subdivided by virtualization software. With blade systems, there may be sharing of input/output resources. With the recent advent of multi-core processor architectures, the processor counts have become a little more confusing, but the following is a good rule of thumb. *If there are many processors and one or a few enclosures, it is “Scale Up”. If there are many enclosures with a smaller number of processors, it is “Scale Out”.*

2. **Policy-driven execution and optimization.** This is, sort of, an “anti-chaos” statement. If you can’t set priorities, you can’t optimize. Larger enterprises operate on such a large scale that they must optimize their IT infrastructure to maximize the ROI of their IT investments and operating budget. Quality-of-Service is the basis for measuring everything.
3. **Dynamic, workload-driven resource allocation.** This means bringing the necessary resources to the work that needs to be done, without much over allocation. This is the opposite of preassigning assets (such as a set of servers) to run a set of applications under a wide range of business conditions and workload assumptions, which almost always results in serious overallocation of resources.
4. **Dynamic workload management.** This means that applications and users get what they need (in terms of resources), when they need it, within policy authorizations. Think of the “workload” as the set of things that have to be done in a given span of time.
5. **High-availability.** For mission-critical applications, this means 24 by 7, year after year without disruption, either planned or unanticipated. Would you ever want it any other way?
6. **Protected execution.** This means that independent work (think “applications”) cannot be accessed by other applications or users, intentionally (say, by hackers) or unintentionally (say, by an address violation or buffer overflow).
7. **Dependent execution.** Some things (think “processes”, but it also applies to data use) need to be done in a specific order.
8. **Guaranteed execution.** Important when something needs to be done within a certain period of time (responsiveness) or at a specific time (scheduling).
9. **Controlled access.** No one (inside or outside) can initiate work or see or change data unless they have the right permissions. A single point of control for security is preferred.
10. **Audit and accounting.** Nothing happens without it being recorded (“logged”) and accounted. Chargeback is optional, but recommended. There is no free lunch.
11. **Infrastructure scalability.** You want to be able to grow your underlying infrastructure transparently, broadly, and efficiently.
12. **Hefty pipes.** You want to be able to keep up with I/O-intensive demands, which typically are found in commercial, transaction-rich applications.
13. **Manageable.** It must cost less to manage (per unit of work) as the collective workload scales larger.
14. **Efficient.** It must cost less to operate (per unit of work), as the collective workload scales larger.
15. **Friendly.** It must play well with other infrastructure assets, whether from the same vendor (or not) and whether genetically similar (or not).
16. **Flexible.** Change just keeps on happening! Unless you intend to be frozen at a point in time, your solution should be conceived and implemented to support what was, what is, what is yet to be. To plan for a success at a future point without retaining flexibility for change will likely end up in failure.
17. **Durable, adaptable, and evolving.** The scaleable architecture should evolve so that wholesale replacement of the infrastructure is not required. A history of successful adaptation to changing requirements and the advances of technologies is a real plus. (It is harder to adapt than start anew; this is the real lesson of “survival of the fittest”.)
18. **Hardware transparency.** This should be a virtual infrastructure of many dimensions, including resource components (processing, memory, storage, network bandwidth, etc.), time (synchronization), and space (geographical distance).
19. **Software transparency.** Linux, Java, and Service-Oriented Architectures (SOA) are the *lingua franca* of the 21st-century data center. What lives beneath should not be important, except for meeting service requirements. This makes possible the transparency of hardware and operating environments.
20. **No lock in.** Applications, processes, and data can run and be stored on infrastructure from a variety of vendors.

These long-standing characteristics seem to make sense for most medium and larger enterprises. None is architecture-specific, but some architectures and approaches may fit better than others. Therefore, you might conclude that I am an *operational IT zealot* – one who appreciates what it takes to run high-volume, complex (interrelated), and unpredictable operational workloads, typically found in larger enterprises.

Are you with me, so far? OK, I have just described *IT nirvana*³. **Many of us covet a solution that can meet these requirements, but only a small fraction of us understand that this can be done better with a Scale-Up architecture and that this world is dominated by today's mainframe.** To further your understanding, let's look at 20 mainframe myths that might explain why this is so hard for so many to believe.

20 Myths to Reconsider ... and Abandon

One of the nice things about labeling something as a *myth* is that the reader knows where the author stands. If it is a myth, it is – by definition - not true, but commonly believed to be true. So, here are many of my favorite mainframe myths, which keep many folks from considering the mainframe path to IT nirvana.

1. **The mainframe is dead, extinct long ago. No one uses them anymore.**

False! Most of the world's largest enterprises rely on mainframes for the bulk of their mission-critical transaction processing and data serving. Far from extinct, this dinosaur has survived and continued to evolve. Server growth for mainframes has outpaced other designs. It is the most sophisticated and most secure IT system available. It is often said that 80% of the world's mission-critical data is stored on and accessed by mainframes.

2. **The mainframe only runs legacy software (like CICS, IMS, COBOL, etc.).**

False! The majority of the new engines (the "processor complex") sold are for Linux and Java applications.

3. **The mainframe makes no sense for an open systems data center (with no legacy applications from the mainframe era).**

False! The majority of the new engines sold are for Java and Linux. That is where most of the enterprise-class application development is being done. If you are running SOA workloads, the mainframe also makes sense, because most executables are Java-based or have components that run in a Linux environment.⁴

4. **The mainframe is unaffordable. It just costs too much to buy and run.**

False. While the hardware appears to cost more on an application processor basis, the mainframe does much more work. In addition, the mainframe has many supplemental (hidden, special-purpose) processors (for I/O, security, and management) that are not included in the classical "engine count"⁵. Furthermore, the fastest-rising cost today is for systems administration, where the mainframe shines as a lower-cost operating environment. Lastly, the mainframe may be the most efficient enterprise-class solution, in terms of electrical consumption, floor space, and cooling.

5. **The mainframe makes no sense for mid-sized enterprises.**

False! The new System z9 BC (business class) is targeted at mid-sized enterprises and offers a broad range of sizing options and upgrade alternatives.

6. **The mainframe requires special skills that are increasingly harder to acquire (i.e., training) or find (i.e., hiring) and afford.**

False! Mainframes require fewer people to act as administrators than what is required to manage so-called "open systems". Training is readily available and many thousands of students are now in training programs around the world.

7. **The mainframe requires expensive disks.**

False! The mainframe has its own storage networking protocols (FICON and ESCON), but there are mid-range-priced storage arrays available, including IBM's System Storage DS 6000, which can be

³ An ideal condition of rest, harmony, stability, or joy. (Dictionary.com) To me, *harmony + stability = joy and rest.*

⁴ If you can define your world solely in terms of Windows-based applications, then you probably are locked into that world. The debate on whether Windows is "open" will not be addressed herein. Or, if you can satisfy all of your enterprise IT requirements with Windows servers, then much of this bulletin may not apply to your organization.

⁵ As in so many other ways, mainframe terminology is different from that of traditional SMP servers. Mainframes contain "engines", while SMP servers run on "processors". There is more than a subtle difference here. The *engine count* in a mainframe refers to the number of IBM System z9 general-purpose and special-purpose processing complexes. Very importantly, there are many "extra" processors in each mainframe, but only the ones used for hosting operating systems or running special-purpose software are counted as engines. This is contrary to most Scale-Up SMP servers, which rely overwhelmingly on its "standard" processors (RISC/Intel/AMD) for all functions.

used as FICON/ESCON and/or Fibre Channel storage. In addition, there are third-party controllers that allow connection of Fibre Channel storage devices to mainframes.

8. The mainframe isn't based on an industry-standard architecture.

False! (This subject is worthy of a couple of pages, but I'll be concise here.) Yes, the mainframe uses its own (unique) processors. Yes, they are different from Intel/AMD, but what difference does that make? If the complaint is about the cost, judge the Total Cost of Ownership. If the complaint is about the instruction set, they all do Linux and Java (which were designed to separate the chip architecture from the application environments). While "Industry Standard" sounds appealing, those processors are less advanced (by a long mile) than the z9 processor. Also, recognize that z9 is an evolutionary child of 40 years of processor design, with backward compatibility along the way.

9. The mainframe isn't Scale Out in design, so it is archaic (i.e., Scale Up is dead).

False! (Again, this is worthy of a longer discussion.) "Scale Up" is better for larger applications running many workloads in a shared-resource environment. Sure, it is easier to think about mono-through quad-processors on a board or a blade, but that doesn't make them superior for commercial processing being done to meet quality-of-service requirements. The reality is that virtualization should hide all of this physicality from developers and users. It then becomes a discussion of how best to match many workload needs to policy requirements. Scale Up does this better!

10. The mainframe has no place in multi-tiered architectures (with application servers, database servers, and web servers).

False! These application architectures run better on a mainframe, especially with the special-purpose engines (for Linux, Java and database) because they can be connected at high speed via HiperLinks (rather than connecting externally through the network). In addition, the mainframe can be used as a data server to applications running on other platforms or can serve as a "hub" that controls the execution of application workloads across a heterogeneous set of servers.

11. A mainframe locks you into IBM for platforms, storage, software, and (many) services.

False! While IBM is the last of the big computer vendors to offer mainframe computers, they are not the only vendor offering "mainframe compatible systems". Additionally, you have many choices for mainframe storage, applications, operations, and management software. Contrary to what is a popular belief, you don't need IBM services to run a mainframe data center.

12. If I don't know about it, the mainframe must not be relevant!

False! What we have here is a "knowledge gap", which is closely aligned with an age gap. Those of us who have been in IT since before the PC era (say, before 1985) generally understand large-systems thinking and its close relative, the enterprise data center. If you are younger than 35 or 40, then you might not have had any significant experience in a "shared computing" (a.k.a. "time sharing") environment. These days, many computer science, information systems, and IT faculty don't have large systems experience. Just because they or you haven't seen it, doesn't mean that it is archaic.

13. We just got rid of our mainframe, right after Y2K. It can't be the right system just five years later.

False! That was your grandfather's mainframe. Time marches on. The dinosaur evolved into a bird...the mainframe now is an eagle.

14. We don't use terminals any more.

True! (Just wanted to see if you were paying attention!) Nobody uses terminals anymore, but "thin clients" of one sort or another, do seem to make sense for many applications. With mainframes, you can use just about any client device (from PCs to web-enabled phones to thin-client stations, no different that most other servers) This is an application issue, not a server issue. However, mainframes will allow access to legacy programs that presume that you are on a terminal, even though you are on a more modern piece of equipment.

15. _____ (fill in the blank) will say that I am nuts, if I consider a mainframe. (Suggested entries: my IT staff, my boss, the CIO, the CFO, etc.)

Well, maybe. But that is not reason enough to ignore this possibility. You and your boss may become heroes, if you can meet your enterprise's many requirements (for quality of service, flexibility, security, etc.) with a mainframe and lower the total cost of ownership (TCO)!

16. I can't get to a mainframe environment from where I am. Too much change and pain. It makes no sense.

False! Spoken like a true coward. You won't know until you check this out.

17. My enterprise's "open systems" applications are not supported on the mainframe.

False! Most leading enterprise applications will run on the mainframe. Most run on Linux platforms or are Java-based!⁶

18. I run Oracle DBMS. The mainframe makes no sense!

False! While IBM would like you to use DB2, Oracle does run on the mainframe.

19. IBM has priced the mainframe as a sucker play. The more I use it, the more costly it will get! Software licensing is the killer!

False! The more you use it, the less it costs per unit of work done. That is why the largest of enterprises love it. Admittedly, getting started from scratch will have higher costs, per unit of work done. However, this is about quality of service and asset utilization, where the mainframe is at the head of the pack.

20. Working on a mainframe will ruin my IT career.

False! Once you look under the hood, you will understand a different meaning of "cool". (More like a turbo-charged "NASCAR" than a standard vehicle at your dealer's showroom. Vroom!) Mainframe experience is a high source of job security and growth. This is about being part of the "ultimate infrastructure".

Conclusion

After 40 years, the mainframe has proven its worth to enterprises, but there is one thing that needs to change, the word "mainframe".⁷ It comes from a reference to an enterprise's *main* or principal computer (for commercial processing), as opposed to the many single-purpose (use-specific) computers of the 1950s. By the 1960s, every large-scale multi-purpose computer was seen as a "mainframe".

While I am comfortable talking about mainframes, many are not, because they just don't have the multi-decade historical perspective. So we need a new moniker, one that speaks less to the physicality of the hardware and more to its role in the enterprise. How about "UltraPlex", to represent the ultimate solution to a complex set of enterprise challenges? Every larger enterprise probably needs one (or maybe a set of heterogeneous resources, with a mainframe at its core)! Whatever you call it, don't overlook your enterprise's need to meet complex, multi-faceted workload (processing) requirements, including access and upkeep of enterprise information collections (databases)! Making all of this happen is where the mainframe has no match.



⁶ Footnote #4 also applies here.

⁷ Many thanks to my colleague John Young for making this so clear to me.

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 cofounder 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 third 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, 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.