THE CLIPPER GROUP
Navigator™
Navigating Information Technology Horizons℠
Published Since 1993 Report #TCG2013019 October 16, 2013

Why More of Your Mission-Critical Business Should Be Processed on IBM System z

Analyst: Mike Kahn

Management Summary

Time and money – these increasingly are among our scarcest commodities. In the 21st Century, there never is enough time. Maybe a decade or two or three ago, we might have felt that time was getting short. Now, there is no doubt. **There always seems to be a need to do more – in less time.** Often we must prioritize what needs to be done first. The result may be that some things don't get done in a timely manner or may never get done. **Money – is there ever enough?** Not only are most budgets tight, but we often face increasing pressures to do more (a lot more) for the same or less than the year before.

Is this about each of us in our personal lives, about each of us in our professional responsibilities, or about the challenges facing our enterprises, especially with respect to information systems? The answer is "yes" to all three. All three perspectives share much in common, as you will see. **Clearly, time and money are critical variables in the formulas of existence and survival in the 21st Century.**

Distance is a third critical variable, in so many ways. Distance often sits on the other side of the equation (from time and money). With a little algebraic manipulation, we can determine the distance traveled in a given amount of time and distance traveled for given cost (i.e., for a given amount of money). Good examples are *miles (kilometers) traveled per hour* and *cost per mile (kilometer) traveled*. These are useful metrics for measuring the efficiency of the distance traveled.

This paper is about time, money, and distance – but not in the abstract. It is about large-scale business processing and management of the related data. **Core business processing – handling the transactions, recordkeeping, data storage and retrieval, etc. – is about time, distance, and money, with an emphasis on time and distance.** *Why focus on time?* Because time may be the biggest challenge that we face because we just can't make more of it. This reality makes time the scarcest commodity. *Why focus on distance?* Because it often correlates directly with time. *Why money?* Because our needs and desires usually exceed our restricted budgets and limited resources.

The big question is: *How can we do more work in a given (limited) amount of time?* In this case, work means *enterprise business*. Narrowing in on this a little, the question becomes: ***How can we do (process) more mission-critical business transactions while still being timely?*** This is both a line-of-business question and an IT question. Doing "more business" is a near-universal business objective, one that usually is subject to limitations involving time, money, and distance.

While this is a broad subject for which many generalities may apply to any platform where systems of record exist, this paper will focus only on IBM System z mainframes running z/OS. Read on to find out why you should focus on System z and how it can process more business transactions – in spite of the limitations of time, money, and distance.

IN THIS ISSUE

- Why Focus on Mission-Critical Business Processes and Transactions? 2
- Why Focus on System z with z/OS? 2
- How and Why Does Distance Matter? ...2
- What about the Money Variable? 5
- Meeting the Objectives on System z with z/OS 5
- Conclusion 8

Why Focus on Mission-Critical Business Processes and Transactions?

In most every enterprise, nothing is more important than *core business processing*. That's why it is labeled as *mission-critical*. To be clear, in this paper "processing" is being used quite broadly. While it once may have been synonymous with tightly-controlled transaction processing (commonly referred to as *OLTP*, short for *online transaction processing*), today it would be better to describe core business processing as the centralized execution of business and the permanent repository for related recordkeeping. While it sounds very similar to what we might have called OLTP, the difference today is that the nature of transactions may be much more broadly defined.

Another way to say this is that the focus is on *systems of record*, which are the engines of execution and recordkeeping for mission-critical business. (See Exhibit 1, at the top of the next column, for a brief overview of the concept of Systems of Record.)

Exhibit 1 — What is a System of Record?

A *system of record* is where business transactions are executed and where the permanent operational and transaction data is stored (usually in databases). This is in contrast to a *system of engagement*, which is where the user interfaces with the business. For example, the system of record may hold the data on what is in a warehouse and which customers have ordered what and when, while the system of engagement might be a web portal that manages customer interactions, whether in-house, through resellers, or from mobile devices. Of course, the system of engagement interacts with the system of record. Today, responsibility for each of these systems often tends to be separated organizationally, for a variety of reasons.

Why Focus on System z with z/OS?

If an enterprise cannot transact business (in the 21st Century, this is presumed to be done electronically), then its ability to survive as a business will be threatened seriously and quickly. **This is the first reason, and possibly the primary reason, for putting mission-critical workloads on *System z* with *z/OS*, as it is the commercial server platform with the highest reliability, availability, and qualities of service delivery.** This is a presumption to what follows. If you already have a Systems z, then you already know this. If you need convincing, you need some basic education on mainframes, for which there are many papers on Clipper's website.[1] There are many more reasons than the one just described. Here are two more.

Second, many if not most of the largest enterprises have entrusted System z with z/OS to be where its most important systems of record applications are run and where their permanent records are stored. Given that this is where much if not most of the "heavy lifting" (of mission-critical application execution) is done, focusing here just makes good sense.

Third, System z with z/OS has a number of features and capabilities that make it the superior platform for mission-critical workloads. (These are the reasons for hosting the most important and time-critical applications on a mainframe.[2]) Many of these features will be discussed in this paper[3].

The purpose of this paper is to present some new thinking to the folks who manage and deliver mission-critical enterprise business processing and recordkeeping on System z with z/OS. It is intended to explain why new thinking may be advantageous. It is not intended to be a deep dive into how to use all of the features and capabilities of System z with z/OS and value-adding IBM software. Please read on.

How and Why Does Distance Matter?

In the Management Summary discussion about time and money, distance was introduced into the equation. Distance is at the heart of our focus on mission-critical systems of record. In our daily lives, we know that distance matters. The concepts of *near* and *far* are intertwined with what we do each day. Intuitively, we all accept that it almost always is quicker to deal with something

[1] For more background on System z, click on the following URL, which will execute a real-time search for "System z" on more than 800 of Clipper's publications on clipper.com: http://search.freefind.com/find.html?oq=mainframe&id=12264876&pageid=r&_charset_=UTF-8&bcd=%C3%B7&scs=1&query=System+z&Find=Search&mode=ALL&search=all

[2] The second and third reasons are intertwined. One could ask: *Which is the cause and which is the effect?* In this paper, they are taken together, sort of as a combined yin and yang.

[3] And many more are discussed in the collection of mainframe papers referenced in footnote #1.

that is near, whether it is a restaurant around the corner or a neighborhood school several blocks away. The importance of distance is magnified (actually, multiplied) by the frequency that you travel the distance. While this has many personal ramifications, **the effects of distance are even more important for a system of record, where the distances may be traveled many thousands or millions or even billions of times per day.**

With transaction volumes going up (which usually is a good thing from a business perspective) and with transactions getting more complicated (often by involving diverse data collections and multiple-step decision and transaction processes), **shortening the distances to be traversed and reducing the number of times that it must be traveled should be considered as attractive ways to accelerate – noticeably – core business processing, primarily because of the very high volumes of transactions in larger enterprises.** Multiply the distances traveled by a million or billion times a day (or more) and the product gets to be very large – very quickly. Even when traveling at the speed of light, the laws of physics are working against you, especially with mega- and giga-transaction volumes.

But it can get worse! On a personal level, if you are driving to work (whatever the distance), if you have to stop at traffic lights repeatedly or get slowed or stalled in traffic, the accumulated delays can greatly increase the time it takes to get to your destination. **In terms of information systems, there potentially are many such delays, such as waiting for data to be fetched or waiting for shared server resources** (e.g., the processing power to execute an analysis or next step in a transaction process) **to get you to the desired result or end point. Multiply that times millions or billions (whether per hour or per day) and you know why optimized systems performance can make a big difference.** Of course, you may be thinking that you could just add more hardware to handle the heavy load, but that usually comes with a heavy financial cost. **It always is far better to do more with what you already have.**

But there is even more to interfere with getting the task or transaction done! This additional factor might be called *The Great Interrupter*, because it isn't measured in milliseconds and microseconds (like most computer-based processes). The Great Interrupter more likely is measured in seconds, minutes, or hours. Humans work at these speeds. **If a business process requires human involvement, e.g., for intervention or approval, that may take a modest amount of time (in human terms) per involvement.** Multiply that times many occurrences per day (dozens, hundreds, thousands, or more), and the speed of business decelerates quite noticeably.

But it can get even worse! Humans are not automatons. Humans only can work so many hours per day and only so many days per month. Obviously, if what they do is complex (like multivariate analysis, such as insurance underwriting or exception approval), work may back up and it might take quite a while (especially in comparison to sub-second information systems speeds) to work through the accumulated backlog. **Often, every request becomes delayed, as each tends to be processed in the order received. Then the delay becomes the norm. This quickly becomes a serious, permanent problem.**

Even if you assume that all systems can "process" at the same speed (a false assumption, but a conceptually simplifying one), the time to travel the required electronic distances and encounter the many delays along the way (both within the information systems and especially the human handling and intervention) will affect how much business can be transacted in a minute, hour, or day. This is at the crux of many of today's business processing challenges.

How and where you choose to execute your business processes will make a difference. Anything that you can do to shorten the total distances traveled (both literally and figuratively) and to reduce the delays and human interventions along the way will make a very big difference at high transaction volumes.

So, what does this mean in practical terms? To answer that, we must turn from the abstract to the technical and look at the details of the underlying information systems technologies.

Addressing the Problems of Time, Distance and Delays – from a Systems Perspective

Let's explore this by an example. If your business process (say, a transaction, such as an order) requires data from multiple sources (think of these as *repositories*) collectively containing a variety of data (like customer information, credit information, prior order and payment histories, current inventory levels by item and location, and demand-driven pricing), it is very likely that the data for completion of this transaction is coming from several repositories. Getting all of the needed data to the application that is processing the

transaction involves many requests, each of which has to travel a distance and may be subject to delays. **Thus, shortening the distances and reducing or removing the potential for delays will shorten the average time to completion.** Doing this on a "faster system" (a server with many high-performance cores with lots of memory and faster storage) will only make this better once all of the needed data is in place and human involvements have been either eliminated or reduced in frequency or duration. **Thus, going faster helps, but not that much in comparison to the other possible solutions for more timely execution. You also must address the distances traveled and delays incurred.**

How to Shorten the Distances

In simple terms, this is about bringing the data closer to where the processing is done. And it may be that simple. **If all of your data is located within your systems of record, where your mission-critical business is processed, then the distances are shortened (in general, when compared to the data being housed and/or processed farther away)**. While the distances can be shortened (and thus accelerated) by the use of certain technologies, such as in-memory indexing, by the use of flash memory and storage devices, and by the use of server busses (and memory-to-memory communications) instead of local external networks, they only will make a big difference when the processing and data are nearby.

So, whether you move the data closer to the processing power (or move the processing power closer to the data, another approach), you need to shorten the distances as much as you can, especially when transaction volumes are very high. The flip side of this is that you want to avoid relying on data and/or processing at a distance (i.e., are far or even farther away), for all of the obvious reasons that will cause more waiting for data or might incur other delays.

While distributed processing and storage (at arm's length or farther from System z) may sound attractive and may appear to be more cost effective (which is debatable, in light of System z's tightly-coupled outboard x86, *Power Systems* and *IDAA Netezza* servers[4]), if multiple points of processing and data storage (even when adjacent) are involved and, thus, all of the needed data and interim decision results need to reach System z before a core business transaction can be completed, then distance will matter. Potential delays, especially for human intervention and decision-making, will only make this worse.

What to Do about Human Involvement and Intervention

Let's assume that all of the relevant data has been assembled and that dependent processing (like credit scoring) has been done and the answers are in. Why might human involvement or intervention (a.k.a. "decision-making") be required? There are several possible answers, some seemingly good and others maybe much less so. Increasingly, this boils down to whether the process can be described as *algorithmic* or not. Algorithmic processes follow a rational, predetermined set of steps (usually involving the comparison of relevant data to established norms). These can be expressed in human terms as procedures, rules, policies, and the like.

If these all are data and rule driven then the process can be done without human involvement (once the rules have been approved), i.e., by automatic execution of the algorithm. There may be times or situations when human review, additional data, and/or approvals may be necessary, but these should be rare exceptions and not the standard practice, for all of the delaying reasons mentioned earlier. To wait for human review of a "mechanical" (i.e., fully-scripted, algorithmic) process just lengthens the process unnecessarily and usually is accompanied with a significant potential for delay.

Your goal should be to automate decision-making wherever and whenever possible and to track all automated and human-intervened processes to identify when, where, and why completion is being delayed. While this may seem like common sense, in retrospect we often find that many of the things that we do (i.e., still do) no longer make sense. These need to be identified both by analysis of processes and by thorough identification and correction of real-time bottlenecks (a.k.a. "hotspots") – so that they can be resolved and automated, whenever possible. Visualization (via information processing and data flow diagrams) provides exceptionally valuable aids to identifying and resolving bottlenecks. Disconnected flows, connections to external systems and storage, and human involvement (and much more) all tend to jump off the page when visually expressed. When business processes and data are documented, kept current, and interactively available, the unresolved inconsistencies and illogical

[4] For more on this subject, see **The Clipper Group Navigator** dated December 12, 2012, entitled *Addressing New Business Analytics Challenges - When the IBM zEnterprise Really Makes Sense*, which is available at http://www.clipper.com/research/TCG2012030.pdf.

or inefficient processes can be assigned for correction or streamlining.

How Can All of This Be Summarized?

In order to do (or be ready to do) more work in the same amount of time (or shorter) in a resource-efficient way, we need to do the following.

1. **Rationalize and simplify (streamline) the mission-critical business processes by (continually) identifying bottlenecks and delays.** Let's call this *tightening the processes and information flows*.
2. **Increase our systems efficiencies by "shortening the distances"**. Let's call this *tightening the IT infrastructure*.
3. **Automate policy- and data-driven decision processes by taking humans out of the loop, hopefully much-to-most of the time.** Let's call this *intelligent automation, whenever possible*.

These tightening and automating objectives can be described in many ways, which can lead to confusion, especially if you are gathering data and processing input from many sources and systems. Our goal is not to add to the vocabulary and confusion but to bridge it and, when necessary, explain it.

What about the Money Variable?

Hopefully, we all agree that we need to simplify the processes, shorten the distances, identify and reduce the delays (whether from information systems components, business processes, or human involvement), all while taking advantage of what System z with z/OS has to offer. Of course, we want to do this in an economically-efficient way, but cost reduction is not our primary goal. As a reminder, our goal is to enable the business, reduce time to completion, and properly handle the data of already large-scale System z-based business transaction processing that tends to be growing in volume, complexity, and importance. In short, we need to do more work without bursting at the seams. Handling more work likely will cost more, which is OK on the assumption that doing more business is desirable and profitable.

A better metric than total cost is the cost per transaction. If we can keep the additional work from raising the average cost per transaction, that would be good. If we can lessen that average cost with the additional volumes of work, that would even be better.

In addition, we want to have and use optimally the systems resources that are needed to make all of this happen – without overdoing it (by overprovisioning wastefully), all while being able to handled greater volumes (potentially, significantly greater volumes) that might occur unexpectedly.

Meeting the Objectives on System z with z/OS

As our attention now turns to systems of record running on System z with z/OS and related IBM software, we need to map our objectives and challenges to what IBM has to offer, since IBM has the richest set of offerings for a System z environment.

IBM has many products and services and presents many concepts to its customers and clients. Let's examine a few and map them onto our three summarized objectives. Be aware that there may not be a single software product that answers all of the needs. It may be an amalgamation of several or more. *It is not our objective to identify, describe, and explain every possible IBM offering or feature that contributes toward achieving the desired goal(s), but to point you to key enabling components for your consideration and review.*

Tightening the processes and information flows.

Often this is called *Straight-Through-Processing*.[5] The word *straight* is key to understanding this. **We really want each kind of business transaction to be processed in as close to a linear fashion as is possible. What this means, in a perfect situation, is that all of the processing is done on System z, that all of the necessary data resides within System z, and that all of the recordkeeping is stored on storage attached to System z.** What is avoided are dependencies to what resides or is processed on other (non-mainframe) servers and storage, especially when many loops or iterations of outboard activity are required to get to that outboard data or processing. By straightening (tightening) the System z-based processes and information flows, everything should go faster, especially when z/OS optimizes the many workloads being executed simultaneously. **What you want to do is to allow System z and z/OS to do what it does best, which is process transactions, with a minimum or no external interaction. If you can do straight-through-processing on System z, it rarely will make**

[5] See Straight-through processing. (2013, March 22). In Wikipedia, The Free Encyclopedia. Retrieved 17:01, October 14, 2013, from http://en.wikipedia.org/w/index.php?title=Straight-through_processing&oldid=546388600

sense to do any part of it on another platform or storage device.

The big question is how to do this. It can be really hard, especially with complex business processes, without a good road map that describes the lay of the land, the way that pieces and places are connected, the vehicles for information flows, the points of decision-making (and possible delays), etc. Let's call this a need for a "process and data map". While you may be thinking of this as something that is done once and remains generally static, that is an outdated thought. While the by-products may look like and be used as documentation, a business process model really is a living organism and one that often is filled with mysteries and missing connectors for how to get from here to there. The roadmap may be contentious, because others may see the flows and processes differently from you, often for good reasons. Thus, the map may take many iterations and negotiations for all involved parties to agree on what reality really is.

Given that our focus is on System z, we need tools and analytical techniques to document what is happening currently on the mainframe. As you will find out, sometimes what actually is happening is not what you expected. That's OK, because exploring the territory is part of the challenge, especially since so many systems have been modified or repurposed over the years or even decades.

Doing business process management is a largely human-driven, manual process. There is no one magic button to push to get a complete picture of reality (from any perspective). Just like Google sends out camera operators to capture what is really out there (for better or worse) on every street, trail, body of water, etc., you need to put on your boots and follow the information and processing trails and points of connectivity. **What is needed is an easy-to-use, yet comprehensive tool to aid in the gathering and integrating of what you find out there.** No doubt, there will be many surprises that you will uncover. You don't necessarily need the tool to run on System z (because it tends to be used for higher-value mission-critical purposes for which there is no alternative). **What you need is a tool to describe what is going on within z/OS.** This doesn't require a tool from IBM, as there are many ways that you can gather, document, share, and update what's going on. However, there is a first-rate IBM product, which also has downstream synergies with other IBM features, capabilities, and software packages.

IBM's multiplatform tool is called *IBM Business Process Manager (IBM BPM)*. It can run on all of IBM's platforms and most operating environments, including z/OS. However, most users use it on a *Windows* desktop or laptop. Its use may be greater on the distributed systems side of your enterprise; be sure to check this out, as there are many synergies from doing business process management across the enterprise.

Usually, different analysts, programmers, designers, architects, and users describe the different parts of the information and process flows that they know and/or for which they are responsible. IBM BPM provides a point-and-click GUI that allows for quick entry of application, information, and flow data. Because it is a tool to be shared across the enterprise, it also is a tool that can be used to come to an agreement of data, processing, and flows among the many organizations of the enterprise. In this role, it can be a tool for mediation (resolution of conflicting opinions) and for documenting changes to be made.

IBM BPM is one of IBM's offerings for documenting and tightening mission-critical processes, information flows, and the like. Rounding out the list are *Blueworks Live* and *Business Monitor*.

- **Blueworks Live is a process modeling tool whose results can be fed into the BPM process designer tools.** It provides a web-based process "capture" vehicle which, because it is web based and so accessible to a broad set of users, provides the means to capture and document business processes collaboratively. Blueworks creates artifacts that prime the process designers providing continuity from the efforts of capturing the process flow to the work of designing improved (tightened) processes that IBM BPM will deploy.
- **Business Monitor provides visibility into real-time, end-to-end business operations, transactions, and processes to help optimize processes and increase efficiency.** It allows the involved parties to observe BPM processes in execution. This is the feedback part of the cycle where processing bottlenecks and hot-spots can be identified based on actual usage. This information can be shared with the process designers for continuous improvement.

Collectively, BPM, Blueworks Live, and Business Monitor enable a dynamic, and probably cyclic, set of engagements and interactions for process capture, documentation, review, and refinement. These three offerings are 21st Century tools that

continue to be improved by IBM, both by internal development and acquisition. If you are not familiar with these, you should check them out, both with IBM and with others within your enterprise that already may be using them.

Tightening IT Infrastructure

This section might have been called "Shortening the Distances", but by now you should understand clearly why that really it is about tightening (and optimizing) the work being done by System z and z/OS. Shortening the distances is just one of the many ways that System z infrastructure utilization can be optimized. Nevertheless, don't lose that thought about shortening the distances, because it is a good way to think about a lot of the tightening. **By taking external (non-mainframe) systems and data out of the loop, the mission-critical workloads should execute much faster, because the distances to be traveled repeatedly are shorter and because Systems z is very, very fast.**

While IBM BPM can be used to document and improve the flow of information and processes within the mainframe environment, this is not what this objective really is about. Many features and capabilities on System z with z/OS are unique to the mainframe. These often are what make the mainframe "hum" and "scream".[6] Humming is a good thing. There are many features and capabilities that let System z with z/OS operate so smoothly and quickly, and they are too numerous to attempt to list and explain herein. However, a few representative features will be mentioned by way of example.

- **Shared Resources with World-Class Isolation** – All of System z's resources (memory, networks, engines (cores)) are available for multiple simultaneous uses. This would never work without secure application and data isolation from top to bottom. For systems of record, that isolation is desired, because many applications (like payroll and health records) must run in isolation from unrelated applications.

  However, there are times when different applications want to share the same data conveniently without making a copy for each use.[7] With System z, and especially via *DB2* and *IMS*, multiple mainframe applications can work on the same database without duplication to create another copy (and, of course, without conflict, except when both applications are trying to write to the same record at the same time – something intrinsically disallowed by design).

  In addition, what is in mainframe memory can be shared (when desired) among several applications, using *HiperSockets*. This allows slower external networking (traffic and distances) to be avoided and replaced with much faster memory access. These examples are just a few of many ways that System z with z/OS can of tighten and accelerate the flow and processing of information.

- **Mainframe Elasticity** – There is no system better at managing mixed workloads on the same server than System z with z/OS. Unlike most other systems architectures, System z with z/OS brings the needed resources to the application, rather than moving the application to additional server resources. System z with z/OS does this dynamically (on the fly), based on pre-defined policy rules. Thus, when more resources are needed for a mission-critical application, less-critical applications are given less resources or the work is deferred until a later time.

  Additionally, more than any other server, System z with z/OS is able to run at near 100% of capacity (without breaking). In fact, because of extra engines (cores) on most mainframes, systems can run at more than 100% capacity (for a short period of time without additional charges) and for longer (with additional charges). The mainframe's ability to run at near capacity, to extend beyond its licensed capacity, and to dynamically deliver resources to mission-critical applications as they are needed are three reasons (among many) for doing most of your mission-critical processing on System z with z/OS.

- **Policy-Driven Automation** – This is discussed in the next section, but it is worth introducing here, because if you have all of the data and decision processes on System z, the straight-through-processing includes and is enhanced by fully-automated decision processing. Without human involvement, decisions and transactions can be processed at the highest speeds, because there is no waiting for human intervention.

[6] By "hum", we mean operate very smoothly and without hiccups and by "scream" we mean go really fast.

[7] Multiple copies are a processing and I/O hog, a currency nightmare, and often require many replicated copies, which can be a waste of additional storage.

Intelligent Automation, Whenever Possible

Taking humans out of the loop, especially when they are not really needed to make hard-to-define-in-advance decisions nearly always makes sense. Managing the decision-making processes, especially when some business process is stalled (i.e., waiting for a human decision), is terribly important. This is where the IBM *Operational Decision Manager (ODM)* comes into play.

ODM provides an elegant development environment, along with dedicated, business user interfaces, for automating and governing frequently occurring, repeatable business decisions across processes and applications. IBM is promoting the use of rule-driven decision making by including ODM capabilities in other software products providing a consistent decision-making technology that permits and encourages reuse of business policy in multiple contexts. This past summer, IBM released *IBM Integration Bus V9 (IIB)*, the follow-on to the *WebSphere Message Broker* product, with ODM included in the package for development and testing by the developers. When the IIB decision-making elements are ready for production, appropriate ODM licensing will be required.

When ODM is combined with the broad use of business process management tools (like IBM BPM) and techniques and includes straight-through-processing (whenever possible) on the elastic mainframe, the enterprise now is ready to optimize its mission-critical business processes that run on your system of record.

Let's Not Forget Money

The mainframe has an undeserved reputation for being expensive. As we often hear, "You get what you pay for." If your systems of record already are running on System z with z/OS, then you already should know that the mainframe does more work, pound for pound, than any other server purporting to be an enterprise system. **This should not be about acquisition cost but about cost per transaction at high volumes in a secure, highly-available, and elastic environment.**

Doing more mission-critical processing on the mainframe should lower your cost per transaction as business process volumes grow.[8] That is what is important and what differentiates System z from all the rest.

Conclusion

Time, money, and distance all are important concepts when focusing on mission-critical business processing. In the long run, it's not about surviving the increasing pace of business (although that also is important); it's about using the accelerated pace to your enterprise's advantage, through rational processes and optimized infrastructure.

Unlike many Clipper papers, this one really does not end on a single note. The purpose of this paper was to get you thinking differently about mission-critical business processes and how you might make things better for your enterprise. Thus, you now may be more at the beginning of a journey, rather than its end. There's a lot to think about. Better start now!

[8] This should be evaluated over several platform generations, because upgrading System z has many economic advantages over other platforms that usually are bought (and paid for again and again) every three-to-four years.

About The Clipper Group, Inc.

The Clipper Group, Inc., now in its twenty-first 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 co-founder 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 20 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 start-up – 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.