



Why the IBM Mainframe is the Right Place for Enterprise Systems of Engagement and Insight

Analyst: Mike Kahn

Management Summary

McDonald's created the *Big Mac*® in 1967, an icon with which almost everyone is familiar. It rises vertically in a stack, which makes it a good metaphor for another icon, the IBM mainframe, born just a little earlier (announced in 1964 and first delivered in 1966). Both have a "special sauce." Both are icons with which many folks may have strong opinions, both good and otherwise. While McDonald's Big Mac (the burger) hasn't changed much from its origins, the IBM mainframe has evolved continuously over the last five decades. With the January 2015 announcement of the next generation of mainframe, the *IBM z13*¹, IBM has created its own "Big MAC". MAC stands for Mobility, Analytics and Cloud.² Read on and I will explain why you just might want an IBM Big MAC and why MAC on any IBM mainframe should be on your menu this year.

Like many of its predecessors, z13 scales to what can be described as another new extreme of gargantuan proportions. The details are many, but with 141 usable cores (out of 168 in the maximum single server configuration) and up to 10 TB of memory, z13 undeniably is a very big server (in terms of capacity; more specifically, the amount of work that it can do each second, which is measured in many billions, of course). That explains the "Big" in my Big MAC play on words. Some enterprises need this scale, so being big (and being able to scale even further by tightly coupling mainframes together) is a really good thing. In fact, every enterprise needs something that is bigger than its current appetite. This provides the room to handle planned growth and the wherewithal to handle unexpectedly large spikes in demand.

But what about "MAC"? Mobility, Analytics and Cloud are the focus of this next generation mainframe and also the focus of this paper. Like its gastronomical cousin, there's more to IBM's Big MAC than the moniker represents. The McDonald's burger is delivered as a vertical stack of ingredients that, when taken together, deliver what many would describe as a multi-sensory experience. The mainframe also has a stack, consisting of hardware, systems software, middleware, and applications that delivers both superior transaction processing and very importantly – at the same time – mobile, analytics and cloud solutions, all in a most secure wrapper. In more formal terms, this is about enterprise-class systems of record, engagement and insight, as will be discussed.

The story of z13 and MAC take place on many planes, but two are entwined: (1) *What it is and how it works* and (2) *What MAC can do and why enterprises will benefit*. Fellow Clipper analyst Steve Bartlett has focused on the former in his paper *The IBM z Systems and the New IBM z13 - Ready to Transform Your Enterprise*³. If you want to know the *whats, whys* and *hows* about the z13 platform, Steve's paper is

Editor's Note: This paper was originally published on 1/13/2015, and has been updated to include announcements made by IBM in February.

¹ In this paper, the short form – z13 – will be used.

² As you will see, the MAC acronym and "Big MAC" are metaphors that I use to explain the mainframe's value. MAC is not another acronym from IBM.

³ See [The Clipper Group Navigator](http://www.clipper.com/research/TCG2015001.pdf) dated January 15, 2015, entitled *The IBM z Systems and the New z13 - Ready to Transform Your Enterprise*, which is available at <http://www.clipper.com/research/TCG2015001.pdf>.

IN THIS ISSUE

| |
|--|
| ➤ Responding to the Winds of Change.... 2 |
| ➤ Doing It All on a Single System Makes Sense 3 |
| ➤ IBM's "Big MAC" Ingredients and More 4 |
| ➤ Conclusion 10 |

the place to start. While the Big MAC story is centered around z13 (the new “king of the mountain”), the mainframe MAC story is broader, as most of it is equally at home on the previous generations of the IBM mainframe, including the zEC12 and z196. If you want to learn more about how the IBM mainframe does MAC, please read on.

Responding to the Winds of Change

This is a mainframe story but not a mainframe primer, so let’s dive into heart of the matter.⁴ Many of the largest transaction processors (like financial institutions, credit card companies, and reservation systems) use IBM’s mainframes as their *system(s) of record*, which is where the transactions of business are done and where the records of that business are kept. That the mainframe can do this better and more securely is an irrefutable fact and the rationale for continuing mainframe use and upgrades. In the past, this was what defined the “mission critical” label, and transaction processing and record keeping continue to be most critical.

In the last three-to-five years, other needs formerly deemed less critical now also are mission critical. First came the *systems of engagement* and then came the *systems of insight*.

- **Systems of Engagement** – this tends to be where the rubber meets the road. By this I mean where users of all sorts (folks within the enterprise, business partners, and customers and prospects) access and engage the enterprise applications and enterprise data.⁵
- **Systems of Insight** – this tends to be where data is analyzed, answers are found, predictions and recommendations are made, and conclusions are drawn. Often running applications and solutions commonly called *business analytics*; this is where enterprise data is used

⁴ If you need an introduction to mainframe thinking and architecture, go to the search box on Clipper’s publication page [<http://www.clipper.com/publications.htm#Catalog>] and enter “mainframe”. You will find many relevant papers on this subject. Steve Bartlett’s z13 paper, referenced in footnote #3 on page 1 of this paper, has the latest details on the new IBM z13 and also contains many references to Clipper’s earlier papers on mainframes.

⁵ The systems that engage the customer typically are different from where the business transactions are processed. For example, if you are doing a lot of browsing then adding only a few items to your “cart”, there may be no reason to engage the system of record until you are ready to place your order. Historically, this separation was deemed to be prudent so as not to overload the system of record. However, these days that may not be the best thing for the enterprise or its users, as will be explained.

to enhance the business’ and the users’ experiences and decision-making processes.⁶

None of these activities are new, but current trends make the use of these applications mission critical today. We can blame much of this on near universal acceptance of mobile smartphone and tablet usage. It is the continuation of a trend that began with desktops and laptops (and terminals before that), whereby today each user now has one or more always-on devices at their fingertips. This has changed access and use patterns from what might have been “controlled” (i.e., access was limited or restricted in many ways) to “uncontrolled” (where the users decide what, where, and how data is accessed and when business is done).

Consider this example. In the past, an investor or bank client might have checked his or her current portfolio or balance maybe once a day (or even less frequently – maybe once or twice a month). Today that might be done many dozens or hundreds (or even thousands) of times per day, many times automatically (say to show the up-to-the-minute values of a portfolio, which by definition is an always-moving target). Thus, while in the past the systems of engagement might have reflected historical data (e.g., based on last night’s or end-of-the-prior-month’s balances), today’s business is being done on current (up-to-date) data. As you will see, this changes a lot of the prevailing assumptions and practices for accessing the latest enterprise data.

But, as they say in those annoying TV infomercials: “Wait. There is more!” Engaging with users whenever and wherever just isn’t sufficient any more. Enterprises need real-time predictive analytics to optimize business and the users’ online experiences. This is where systems of insight come into play.

There’s no doubt that enterprises need all three and that today all three are, each in their own way, mission critical. Historically, systems of engagement and insight have been loosely linked to enterprise data in the system of record, usually by exporting the data from the system of record to one or more other platforms, where it was cleaned, prepared, and made available to deliver meaningful engagement and insight. It turns out that this relocation and remote processing can be terribly inefficient and costly – and potentially inadequate.

⁶ Historically, these also have been segregated (prevented) from running on the system of record, even though that often is where the best data resides on which to do analytics. This also may be changing, as will be explained.

Understanding a Very Real Problem

It takes a lot of processing time, memory, and network bandwidth (not to mention absolute time) to extract the data from the system of record and push it over the network to other computer systems. So even before the exported data is made available for engagement or insight, it has cost a lot of money and time, during which the data often becomes stale. And that is only half of the story. All of that data that has been exported has to be stored, transformed into something usable and then stored again, possibly many times. This is a never-ending cycle of costs and time – all just to end up with data of varying degrees of staleness.

The obvious question is *why not operate on the data at the point where it is permanently recorded, i.e., run the systems of engagement and insight on the always-up-to-date data sitting in the system of record?* There have been many reasons not to consider doing this, especially fear of somehow affecting or contaminating the processing of mission-critical business and expending the seemingly most expensive computing resources to do what seems to be straight-forward work that could be done on less-capable servers, but none of these needs to be true today.

First, as mentioned before, the costs of exporting, transforming, loading (ETL), and (don't forget) storing and backing up all of the data needed for meaningful engagement and analyses adds up to a big pile of money and the almost always – a lot of elapsed time – before anyone is engaged or any data is analyzed. That will take even more resources. Using your system(s) of record to extract from enterprise databases and then to pump it to other servers likely is wasting significantly more resources than originally was feared to be wasted by actually running engagement and insight solutions on the system(s) of record.⁷

Second, we've come a long way from needing to isolate and protect production databases from other simultaneous uses, especially in mainframe environments. In fact, this is more natural on mainframes than other platforms, as you will see.

Third, and probably most important, is the advantage of having only the most current data available for real-time engagement and analytics. Let me give you a personal example. I was interested in buying a certain electronics

product. Online, I looked at various likely suppliers' websites, including some with local retail establishments. The more I looked, the more the establishments became aware of what I was seeking. How do I know this? Because when I went to an unrelated website for something else, I was getting ads for my sought-after electronics product and related add-ons and alternatives. Privacy issues aside, let's say that this was very impressive, at first, both in terms of engagement and the analytics being done behind the scenes to connect all of the crumbs that I left behind. But it all went terribly wrong and continued to do so for some time. *How?* Well, I bought the product from one of the suppliers. Yet, for more than a month, that same supplier was sending me frequent emails and displaying more ads when I visited its site (and others) – all trying to convince me that I should buy the product that I already had purchased – *from them!* Their systems of engagement and analytics clearly were disconnected from the transaction data of what I had bought from them. They were wasting my time and their systems resources because of this decoupling. **The moral of this story is that stale data usually doesn't generate good business results.**

Doing It All on a Single System Makes Sense

Accept, for the moment, that doing transactions, engagement, and insight on one system just might make good business sense in terms of satisfying enterprise operating requirements and service level objectives. What I am suggesting is to assume that it is technically feasible. That would leave just two important questions.

1. *How could this be done?*
2. *What would this cost?*

Let's take one at a time.

How to Do It on a Single System

Technically, you need a database manager capable of handling many concurrent requests against the same database without compromising the qualities of service (QOS) for any of the applications or users using that database. First, this typically requires some-to-much of the data to be held in memory (think RAM and/or flash) and for I/O to be optimized via hardware and software. This isn't easy and was the primary reason that production work (transaction processing) was isolated from analytical work (especially the proverbial unintentional "runaway query" that has been a source of worry for decades.)

Then you need enough processing power to

⁷ Go ask your data center manager how much of the transaction servers' processing time is being spent on exporting and transferring data.

get it done. While that might seem simple to determine, in reality it isn't, because much-to-most of today's enterprise workload is variable in nature. You might try to over-provision enough processing power (for the moment, simply think about cores and memory) to handle double or triple the anticipated workload (you might call this "just-in-case oversizing"), but that might result in a lot of wasted expensive resources sitting idle waiting for a tsunami to unexpectedly happen. What really is needed is a way to prioritize work so that the resources are appropriately deployed doing something else, until the "big one" happens, when those now-needed resources automatically will be redirected at the most mission-critical workloads and, if necessary, additional resources can be added.

How to Manage the Cost

I could write 10 pages of this, but I will be much more concise on this critical matter. **It often is less expensive to do it all on a mainframe, assuming that you already have or want to have your transactions processed and mission-critical data stored on a secure mainframe.** The reasons are many, including not having to extract it from the database where it normally resides, transport it (by FTP or otherwise) to another server or many servers, load it into an engagement or analytics program, and potentially make multiple copies of the data.

IBM has data and customer experience showing that it costs less to do analytics on a mainframe (and that is without considering all of the costs for multiple copies of the same data that tend to accompany distributed analytics solutions.) **Avoiding the added costs of multiple copies of data (and the software licenses to administer and manage it all) may alone justify the whole mainframe-based operation but you should do your own detailed analysis.** Don't forget to look at middleware and analytics software savings, as many software licenses are based on the number of servers, processors/cores, or, most importantly, the amount of data being analyzed. If you have multiple extracts of the same data or overlapping subsets of the data, then you likely have many copies of the data and may be charged (many) software license fees for all of that duplicated (potentially many times) data – plus the storage capacity it occupies.

While I have just focused on analytics, there are similar stories for mobile users and cloud deployments. In a world of sharing enterprise data, the need to give mobile users secured access to it is not different than getting to it by other

means. What are needed are the necessary interfaces and programmatic gateways, and the mainframe has these as well.⁸

Cloud solutions are similar. If you want to provide secured, shared access to a slice of computing power, provisioned local storage, and access to enterprise data, all of that can be done on a mainframe. *Linux* images and *Java*-focused solutions are supported in spades on specialized mainframe processors⁹ – the *IPL* (for *Linux* for *z* Systems) and *zIIP* (integrated information processor)¹⁰, which not only are optimized for these tasks but also cost less than standard (*z/OS*) mainframe cores¹¹.

In reality, this all boils down to having secure yet shared access by mobile apps, analytic apps, and cloud-based apps to enterprise data. Being close makes a big difference. Distance equates to time and money.¹² That's why you may need a Big MAC (*z13* with all of its possibilities and out-board specialty engines, including *IBM DB2 Analytics Accelerator*¹³ and *zBX*¹⁴) or at least a MAC-focused later-generation mainframe. MAC solutions have been the focus of mainframes for at least three generations (from *z11* to *z13*).

IBM's "Big MAC" Ingredients and More

Earlier, I suggested that IBM's Big MAC had a vertical stack not dissimilar to the infamous burger. I also suggested that the operating system

⁸ For more on IBM's System *z* mobile offerings, see [The Clipper Group Navigator](#) dated October 9, 2014, entitled *When a Secure Mobile Host Platform is Required - Go with IBM zEnterprise*, which is available at <http://www.clipper.com/research/TCG2014018.pdf> and also [The Clipper Group Navigator](#) dated April 23, 2013, entitled *Enabling Your Mainframe Data on Mobile Devices*, which is available at <http://www.clipper.com/research/TCG2013008.pdf>.

⁹ Also called "engines". They are the mainframe "cores" in non-mainframe parlance.

¹⁰ The *zIIP* can also host *zAAP* (application assist processor) functionality, which can be used to offload *Java* execution.

¹¹ In fact, new workloads on a mainframe tend to be priced to compete favorably against distributed Intel *x86* servers.

¹² For more on why distance matters, see the issue of [Clipper Notes](#) entitled *Bringing "the Power" Closer to Home - Implications for Future Storage Architectures*, dated November 8, 2012, which is available at <http://www.clipper.com/research/TCG2012026.pdf>.

¹³ For more on analytics on a mainframe, see [The Clipper Group Navigator](#) dated December 21, 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>.

¹⁴ *zBX* is a tightly-coupled mainframe-connected *IBM BladeCenter* server with Intel (*System x*) and *Power Systems* blades. It allows off-mainframe processing of mainframe data, delivered from a mainframe perspective.

was IBM's "special sauce". Let's dissect the stack, starting with the special sauce.

Operating System(s)

While the infamous burger has a singular special sauce, the IBM mainframe has many. This gives the enterprise the ability to choose what best meets its needs and this likely is more than one operating system. In terms of production capacity, the two dominant mainframe operating systems are z/OS and Linux. Each has its own merits. In some ways, these two are opposites, as I now will explain.

z/OS is the most secure commercial operating system. It achieves this in many ways, but largely by tightly integrating with the mainframe's very secure hardware, which has been at the core of its existence since the beginning more than 50 years ago.¹⁵ This tight integration can be seen as a "plus" or a "minus", based on your opinions about "open systems" versus "proprietary systems". In spite of many features of openness, z/OS is a proprietary system that runs today only on IBM mainframes. It delivers a clearly superior and most secure enterprise solution. If you can live with that or, more likely, that is exactly what you seek, then z/OS is the enterprise operating system of choice.

However, if being open is the driving force in your decisions, then *Linux on z Systems* is the answer for you. You get many-if-not-most of the benefits of running on a mainframe and most of your Linux applications should run without modification. Worth noting, over 7500 IBM and ISV applications and tools have been certified for Linux on z Systems.

There are other mainframe operating systems (including z/VM, z/VSE and z/TPF), each with its own special flavor, and the mainframe can run all or several of these simultaneously. Additionally, another sauce is coming. Based upon its recent statement of direction¹⁶, IBM will be adding KVM to this collection, to give Linux administrators another choice of hypervisor beyond z/VM. **Regardless of whether you want one or several**

¹⁵ For a retrospective on the conception and launch of the S/360 and also why the original design and architecture were right on target for mission-critical business systems, see *The Clipper Group Captain's Log* entitled *The Beginning of I.T. Civilization - IBM's System/360 Mainframe*, dated March 30, 2004, and available at <http://www.clipper.com/research/TCG2004028.pdf>.

¹⁶ Statements of Direction (SODs) are IBM's non-binding announcements of what likely will be released within the next six months. IBM makes these statements to allow buyers to make good, forward-looking decisions that might be influenced by what is expected to happen before very long.

of the special sauces, the IBM mainframe gives you many choices and lets you have more than one.

Having It Your Way

Unlike the burger, IBM's mainframe stack is full of other choices as well. In fact, if you don't mind an altering of a slogan taken from one of Burger King's commercials, it can be said "With the IBM mainframe, you can have it your way!" Regardless of your needs, procuring a mainframe is a custom (made-to-order) experience. You get what you need today and, unlike the burger, you can add to it and improve it as your requirements evolve over time.

What results can be a set of many choices that you have to make when ordering your mainframe. IBM has made this simpler with many predefined solution sets, which can be amalgamated into what you need to order. If you prefer the simplicity of all of your servers coming "just one way", the IBM mainframe may not be for you. But if you want a system that can run and optimize a collection of your solutions, then the IBM mainframe may be the best answer.

What follows is further discussion of some of the key software requirements for your enterprise data, for MAC uses, and for enterprise applications developers. Some of IBM's recent announcements of new capabilities and future possibilities will be presented and discussed.¹⁷

The Exceptional Juggler of Data

Of course, the newly mission-critical workloads today often are of the mobile, analytics and cloud varieties, running against your enterprise databases held in your systems of record (and often now becoming intertwined with them.) Before we get to the MAC solutions, let's first consider what likely is the backbone of your mainframe's enterprise data – your database management software, which in most cases is some combination of DB2 and IMS, although VSAM also is very common for transaction processing.

If anything, **DB2 is the magic that makes a lot of very difficult things possible, including concurrent access to a single database by many users and applications**, including those doing business transactions.¹⁸ Doing all of this

¹⁷ More mainframe software announcements are expected at InterConnect next month. Many of these are expected to round out IBM's "MAC" offering, so do check back next month (at <http://www.clipper.com/publications.htm#Catalog>) for more on the additional announcements.

¹⁸ Of course, this also is made possible by the mainframe platform that allows all of this to be done in a meaningful

simultaneously and without dropping a ball requires a very special kind of juggler, in this case, one that works primarily behind the scenes. Few end users ever knowingly interact with DB2 (there are exceptions), but almost all of us benefit regularly from its exceptional juggling.

Of course, DB2 and IMS (and other mainframe methods of record management) are needed to delivery transaction data (into and out of) systems of engagement and insight application co-located on the mainframe. Thus, we need to build on this co-location and look at some of the many MAC solutions that IBM has to offer, with an emphasis on what capabilities have been added recently.

“M” is for Mobile

Mobile access likely is the tail that is wagging the enterprise systems dog. If the mobile endpoints didn’t exist or weren’t being used heavily during most of the hours of each day, maybe you could get by with a less heavy-duty solution than provided by an IBM mainframe. However, that water passed under the bridge a long time ago for most enterprises. There is no stopping the mobile access phenomenon.

The other way to state this is that enterprises must accommodate mobile users’ needs or face dire business consequences. As we have just seen during this past holiday shopping season, mobile access is driving much if not most of the retail business, either directly or indirectly. Sure, a lot of online business is being done from mobile devices, but that is only part of the story, as business often is influenced by generally-mobile comparison shoppers. It seems safe to say that if you were spending more than \$100 at a retail store (or online), you either comparison-shopped before you reached the store (storefront) or, more likely, you checked competitive pricing and alternative products on your mobile device while you were in the store.

This makes mobile, for many if not most of today’s purchases, whether B2C (business-to-consumer) or B2B (business-to-business), a critical success factor. Simply put, you need a system of engagement for all of these mobile users – or else you risk a declining future. As discussed earlier, this means that you need to tightly and securely couple your systems of record with your systems of engagement.

Last year, IBM released many products to make mobile connectivity to data and business

amount of time and at a reasonable cost for the value delivered.

applications (processes) on z Systems a secure and productive experience¹⁹ and, more recently, announced its first iPhone apps stemming from its joint development partnership with Apple. All of that continues into 2015, albeit with some of the products having been renamed to match the new “z Systems” nomenclature. So, while there may not be a long list of new mainframe “mobile engagement functionality” to accompany the January announcement of z13, do not lose sight that it is mobile apps and mobile users that are driving the need for enterprises to beef up their mainframe footprint to take advantage of the mobile development software was announced last year for the mainframe.²⁰

“A” is for Analytics

As may now be coming into clearer focus, none of this is done in a vacuum. Often it is analytics that the mobile users need to get to what they desire (to find or to do) or it is analytics that enterprises want to use to enrich the users’ experiences and, of course, increase their spending and relationships with them. **So, just as systems of engagement need to have secure access to systems of record databases, systems of insight need the same, but on a much larger scale.**

For example, if I am looking to buy an identifiable product, say a particular brand and model of tablet, flat-screen TV, or power tool, then I am looking for something that is very specific and requires access to a limited amount of data – the data about this single product. Maybe I am getting a modest number of “hits” because the product is available from a number of sources or because it has been reviewed many times. Regardless, this number of hits is small in the universe of thousands or millions or billions of possibilities in an Internet search. However, if I am an enterprise or advertiser wanting to influence the prospective buyer’s decision-making (to buy what I have to sell), I might have to process (analyze) many thousands to millions of records (very quickly) to give the seeker timely and meaningful guidance and encouragement. This requires even tighter coupling with enterprise data, because the analytics engines need to do a lot of crunching in short order for each of the prospective buyers (or other kinds of users) and this needs to be done without affecting the processing of business transactions or other mission-critical applications.

¹⁹ See the papers cited in footnote #5.

²⁰ And look for additional announcements at InterConnect next month.

Newly Available for Mainframe Analytics

- **IBM DB2 Analytics Accelerator** accelerates in-database transformations and analytics. The *Accelerator*, as it often is called, is a logical extension of DB2 that does not require any changes to applications programs, and essentially is transparent to the user. Over the last several years, IBM has built a comprehensive and complete analytics solution on z Systems. Thus, enterprises now can take advantage of their enterprise data where the transactions originate. The DB2 Analytics Accelerator complements the DB2 for z/OS industry-leading transactional processing capabilities, providing a specialized access path for data intensive queries. This enables real and near real-time analytics processing, which executes transparently to applications and users. The Accelerator operates as an integral part of DB2 for z/OS and z Systems.
- **IBM DB2 BLU Acceleration**, previously available on non-mainframe platforms, has been improved and now also is available on Linux on z Systems. It enables faster queries on Linux data marts without time-consuming custom optimization by compressing row-oriented tables and then operating on the compressed data. Because the database is compressed significantly, it takes up less space and thus more data can be stored in the same amount of memory. This is best applied to databases with certain characteristics, including a maximum size of about 10 terabytes.
- **IBM InfoSphere z Systems Connector for Hadoop** enables efficient sharing of mainframe data with **IBM InfoSphere BigInsights**, running either on mainframe Linux for z Systems partitions or on external Intel (x86 architecture) or IBM *Power Systems* clusters. **When IBM InfoSphere z Systems Connector for Hadoop and IBM InfoSphere BigInsights both are installed on the mainframe, the z Systems platform behaves largely as a private cloud.** Hadoop cluster nodes are virtualized, and all operations take place within the security perimeter of the mainframe, simplifying the environment, and greatly extending the range of data that is accessible to the mainframe. IBM InfoSphere BigInsights clusters on the mainframe can run across one or multiple IFL processors²¹,

providing the needed horsepower to make many things analytic happen in real time.

“C” is for Cloud

It has become very “cloudy” just about everywhere. As a metaphor, many of the things that we do each day are being done (or are described as if they are being done) in a cloud of one sort or another or maybe it is just being presumed to be done in this way. I’d even be willing to say that there are so many clouds that it is hard to see through the clouds to what is on the other side – and there always is something on the other side. This is akin to the Wizard in the *Wizard of Oz* and his chiding “don’t look behind the curtain”. **In the real world, whatever is behind the IT curtain tends to be very important.**

Of course, as IT professionals, most of us live on the other side of the curtain (a.k.a. cloud). We know that what seems ethereal (or virtualized into invisible and incomprehensible existence) needs to be real and that every cloud needs to be hosted and managed and secured, etc.

From an enterprise perspective, the cloud usually is partly metaphor and partly real. Put in other words, some of the time it may be provided by another (think distant supplier) without much concern for what’s behind the curtain (except for its ease of management and costs) and the rest of the time we (the enterprise IT folks) are the ones who must make the platforms and services available to their user communities. Many folks speak about public clouds, private clouds and hybrid clouds. You may need all three but **there are many reasons that you need to focus most intensely on your private and hybrid clouds – where your enterprise keeps its most important data.**

First, if you are responsible for enterprise data or application hosting, you need to protect it, whether it is in a cloud or not. Obviously, it is easier to protect it from within (i.e., a private cloud) than if there are public-cloud exposure(s). What you want is a good way to do this on your multipurpose mainframe – where there is no safer place to put up your cloud – which, coincidentally, is where your enterprise data already works and lives.

Second, it seemingly is easy (but not necessarily inexpensive) to let someone else worry about providing the horsepower to slog through your variable and often unpredictable or unknown demands for processing or data. Assuming that all goes well and all requirements are met all of the time, then maybe the public cloud is right for you. Most of the time, it isn’t as

²¹ Or on a standard z processor, but this is not common for production uses.

easy or simple as you might have hoped.

However, if you have to worry about what data is in the cloud, who should be accessing it, what they are allowed to do with it, etc., then you should start to worry about putting your stuff at a distance far on the other side of the curtain. Additionally, if you choose to keep large amounts of current (up-to-date) data in the remote cloud (having been exported from your enterprise systems of record), then you really have to worry about currency (or, more accurately, the degree of staleness), plus the cost of moving the data to that distant cloud.

Newly Available for Mainframe Cloud

- **IBM Cloud Manager with OpenStack v4.2** is an easy-to-deploy, simple-to-use cloud management offering that is built on the OpenStack Juno release to deliver multi-region hybrid support for mainframe-based clouds and is backed by IBM Support. It runs under Linux for z Systems. Multiple z Systems-enabled infrastructure regions can be viewed from a single pane of glass, as can Power Systems and x86 servers, delivering central management across multiple hypervisors and domains. All IBM server architectures and major hypervisors are supported. On z Systems, it also includes support for Chef-based workload deployment based on OpenStack Heat pattern engine
- **Custom Patterns for Linux on z Systems** will give mainframe Linux developers the same kinds of patterns that they have grown to like and adopt on *Power Systems* and *System x* platforms. This allows a pre-configured and pre-tested applications or solutions to be installed with little effort and in short order. Because the patterns are pretested, you are assured that everything that is needed for the Linux application or solution will be provisioned and installed properly. This saves time, effort and improves quality. In phase 1, available now, about 50% of the software portfolio will be available as patterns²². Phase 2 (later) should bring this to about 85% coverage.²³ These new patterns are the result of a partnership with IBM Rapid Deployment Services, which is a part of IBM GTS, but will be available as traditional SWG part numbers that can

²² Calculated as a percentage of mainframe license software revenue to IBM.

²³ Your coverage percentages will vary depending on your enterprise's mix of applications and how you are keeping score (by MIPS consumed, dollars spent, etc.).

be included in enterprise agreements, such as ELAs.

New Software for z Systems Developers

Available Now

- **Java Version V8 Technology Edition for z/OS** is a 64-bit SDK that is designed to be compliant with the *Java Standard Edition 8 (Java SE 8)* APIs. It builds on the "write once, run anywhere" Java paradigm and leverages *zAAP processors* to run eligible Java work which also can be deployed as *zAAP-on-zIIP*²⁴. It also allows applications to more fully exploit the expanded z13 and zEC12 instruction sets.
- **IBM CPLEX Optimizer for z/OS** delivers mathematical optimization on mainframe computers, providing flexible, high-performance mathematical programming solvers. It enables large enterprises with operational data on mainframes to use existing investments and take advantage of the mainframe's processing power and security features. This especially provides optimization opportunities for finance and banking, healthcare, and governmental organizations. The new release will provide support for SIMD²⁵ on z13, which should provide significant performance improvements for numerically intense analytics workloads.
- **New compilers for z13** take advantage of a number of new z Systems hardware instructions and give further-optimized performance when compared to compilers on earlier generation mainframes. These include:
 - **Enterprise COBOL for z/OS, V5.2**²⁶
 - **Enterprise PL/I for z/OS, V4.5,**
 - **XL C/C++ for Linux on z Systems, V1.1**
 - **XL C/C++ for z/OS, V2.1.1**
- **UrbanCode 6.1.1.1** offerings for z Systems allow developers to efficiently and more rapidly deliver mobile, cloud, big data, and traditional applications across hybrid server environments. The offerings automate the deployment of applications, databases, and configurations across heterogeneous development, test, and production environments, and

²⁴ Going forward, zAAP functionality will be delivered on zIIP processors and the standalone zAAP processor no longer will be offered. In other words, you can use a zIIP to accelerate both database queries and Java workloads.

²⁵ Single Instruction Multiple Data, a new parallel execution feature available exclusively on z13.

²⁶ Testing by IBM indicates up to a 14% improvement for COBOL on z13.

helps organizations drive down cost and speed time to market with reduced risk. New capabilities include:

- Automated deployment of application environments to hybrid cloud providers.
- Enhanced automated deployment of applications to enterprise systems.
- Enhanced release tracking to enable coordinated and integrated release management for large enterprises.
- **Rational Test Workbench V8.7** is a comprehensive functional, regression, load, and integration testing solution for all types of applications, including mobile. The new release will support *IPIC* (which enables communication with mainframes using TCP/IP protocols), virtualization, and PL/I. This allows building of intelligent and interconnected applications that can be deployed on traditional and cloud infrastructures.
- **BPM and Monitor on Linux on z Systems** provides tooling and run time for process design, execution, monitoring and optimization of business processes. This is a serviceability enhancement with support for the upgraded Linux operating system.
- **WebSphere Application Server Liberty Profile** has been enhanced with *Java EE7* support for *JavaServer Pages (JSP) 2.3*, *Java Database Connectivity (JDBC) 4.1*, and for the *Java API for WebSocket 1.1* (and more). In addition, there is enhanced *z/OS Connect* discovery service, which improves API management integration with *z/OS* assets; improved *Liberty on z/OS* security integration with the *z/OS System Authorization Facility (SAF)*; enhance *WebSphere Application Server Developer Tools (WDT)* for remote server support, and enhanced security for *Single Sign ON (SSO)* authentication with *SPNEGO* tokens.
- **TXSeries for Multiplatforms V8.2 on AIX and Linux** – Enhanced application development experience, including alternate COBOL compiler support on Linux with *COBOL-IT*, support for *Visual COBOL*, and simultaneous debugging support of TXSeries applications using IBM COBOL. Optimization and serviceability enhancements are for AIX and Linux platforms. This does not reside on the mainframe but interfaces with Linux on z Systems.
- **IBM Integration Bus V10 for z/OS** includes enhancements to the developer experience, improved manageability via web browser

plugin, and cloud support including Chef scripts for simplified provisioning and the ability to publish IIB²⁷ services to cloud integration.

New z Systems Security Offerings Enhance End-to-End Security

- **IBM Security zSecure Manager for RACF z/VM** adds new support for compliance framework and z/VM currency, thus delivering proven-effective RACF administration, reporting, and auditing.
- **IBM Security zSecure Audit SSE** provides effective security compliance checking to *IBM MQ for z/OS*.

Preview of What's to Come (by Recent Statements of Direction)

- **Cloud Integration for Bluemix** will allow mainframe developers to choose sample code or services from a catalog. This will provide a big head start for mainframe application and systems services development and also will serve as a learning vehicle by offering “how to” examples. Expect additional Bluemix services for mainframes in the coming months.
- **CICS Transaction Server V5.3 Open Beta for z13** – Now in an open beta, *IBM CICS Transaction Server for z/OS (CICS TS)* supports 64-bit SDK for z/OS, Java Technology Edition, Version 8 (Java 8). IBM has issued an SOD for the current software to be able to take advantage of the new z13 facilities that can be exploited by Java 8, including Single Instruction Multiple Data (SIMD) instructions for vector operations and simultaneous multi-threading (SMT) on zIIPs.

Conclusion

In this report, I have covered a lot of territory to explain why you want to put your systems of engagement and insight very close to your systems of record and exactly how and why this is especially true for enterprise data centered on the IBM mainframe. While some of the described MAC (mobile, analytics and cloud) software capabilities are unique to z13, most of it also applies to prior generation z11 and z12 mainframes. Whether you are ready for a new (bigger and better) mainframe or not, you should be addressing what it takes to meet your many needs that often are delivered as a result of very

²⁷ IBM Integration Bus (formerly WebSphere Message Broker) is an enterprise service bus (ESB) providing connectivity and universal data transformation for service-oriented architecture (SOA) and non-SOA environments.

important enterprise MAC requirements.

Repeating and slightly adjusting the pun stated earlier, you really need to “do this your way”, that is, delivering enterprise data, applications, and solutions in the best, most practical, most efficient, most cost-effective way that you can.

- If your enterprise depends on a mainframe, then this means doing more on a mainframe.
- If you have you supersized, mission-critical enterprise requirements but do not today use a mainframe, then most certainly you need to look at how you might satisfy your transaction processing and MAC requirements on a mainframe with IBM software.

Regardless, looking at how to do your systems of record, engagement, and insight all on one system may be your most critical mission for 2015. Better get to it before the new year slips away!



About The Clipper Group, Inc.

The Clipper Group, Inc., now beginning its twenty-third 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 40 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 22 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**, **The Clipper Group Calculator**, 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.