THE CLIPPER GROUP
Captain's Log™
Navigating Information Technology Horizons℠
Published Since 2001 Report #TCG2014009 April 11, 2014

The History of the Enterprise Cloud as Seen by an "Old Mainframer"

Analyst: Stephen D. Bartlett

Introduction

An interesting fact (to me anyway) is that the *IBM System/360 (S/360)* is just one month older than the *Ford Mustang*; the S/360 family was announced on April 7, 1964, and the Mustang in May, both in New York City, by the way. In many ways, the product life of this much-coveted vehicle reminds me of and in many ways parallels the S/360 and the 50-year legacy that it spawned. Both demonstrate longevity, certainly, but these remarkable products always have been changing and improving, and have survived many crises, economic downturns, governmental intrusions, technology revolutions, and multiple threats from many competitive venues. Each has endured because of a consistent corporate vision, the dedication of many talented scientists and engineers, and most of all because of the continued enthusiasm, trust, and confidence of its many customers. Even after fifty years, one can never mistake a Mustang for anything but a Mustang. And so it also is with the System/360 and its many successor generations of IBM mainframes. Compared to the mean age of the current flock of IT professionals, fifty years is beyond memory; it may even predate the birth of their parents. My career reference point was set by the announcement by IBM of the System/360, marking its fiftieth anniversary this week. Looking backward through my 50-year lens on the history of modern computing, it seems like a very short time indeed, with the pace of advancement continuing to accelerate.

IBM certainly did not invent the computer nor was the S/360 the first commercially successful computer or the first "mainframe" for that matter.[1] However, as a benchmark of modern commercial computing, there is little dispute that IBM's announcement of the S/360 in 1964 met the burning needs of the world's businesses and public institutions, large and small, in a most extraordinary way. The values that initially made the S/360 so important, so pervasive, and so successful after it was announced – *scalability, flexibility, security, resilience, and cost-effectiveness* – are the very same as those characteristics that are valued and important today – and will be in the future**.** In modern terms, these are the characteristics of enterprise-class computing and, therefore, also are characteristics of enterprise-class cloud computing. If your tenure in the IT industry is less than a decade or two, you need to realize that almost all credit or debit card transaction anywhere in the world eventually will be processed by a mainframe (or, more than likely, several); almost every airline reservation is eventually processed by a mainframe; and almost every claim payment made by large-scale insurers is at some point processed by a mainframe. It is very likely that every person living in any developed nation, and a good many others, has an interaction every day with a mainframe, usually without their knowledge.

> ***Publisher's Note:*** *This is a longer and more personal version of a paper entitled* ***IBM zEnterprise is Enterprise Cloud Infrastructure****, also written by Clipper analyst Steve Bartlett, and available at* www.clipper.com/research/TCG2014008.pdf*. In this longer version, Steve also reflects on his mainframe career at IBM and Clipper, which encompassed almost all of the mainframe's first five decades.*

[1] The use of the term "mainframe" within this paper refers to large-scale (relatively speaking), centralized computing and data processing systems typically found in many larger enterprises and government agencies. They most frequently are used as *systems of record* and serve as the primary processor of secure, mission-critical transactions and as the secure repository of important business data. Its antecedents are the big room-sized boxes enclosed behind glass walls. (For a photo of an early IBM mainframe, see Exhibit 1, on page 3.) Unless otherwise stated, all uses of "mainframe" in this paper refer to IBM's mainframe technology. (See also Exhibit 2 on page 4, which described most of the other mainframe players of the 1960s and early 1970s.)

While there are many historical references and stories in what follows and most of this was brought into focus by the fiftieth anniversary of the IBM mainframe, I contend that much present-day cloud thinking can be traced to many of the mainframe's stepping stones set in the past**.** To this end, this paper has three intertwined themes.

- The history of the IBM mainframe over its first 50 years.
- The concept of "cloud" which, while relatively modern terminology, has been part of the mainframe's persona for more than four decades.
- My mainframe career with IBM and Clipper, which has paralleled these 50 years.

The journey of IBM's mainframe systems will obviously continue beyond these 50-year milestones and my career. History tells us that it's form and function will most certainly change, perhaps in unpredictable ways, but always to remain relevant to its clients and to the marketplace at large. Read on for a true and, I hope, interesting trip through the last 50 years.

Milestones of Mainframe Computing

If one were to list all of the innovations of the IBM mainframe over the last 50 years, that list would have many hundreds to several thousand items. Instead, I have chosen seven significant technological milestones from the history of IBM's enterprise-class mainframe computer to use as a backdrop to make my case about today's mainframe as superior computing and delivery platform for enterprise-class cloud deployment.

(1) **The introduction of "general purpose" systems**, which began with the announcement of System/360 in 1964.
(2) **Large-scale virtualization** beginning in 1965 with the *System 360 Model 67* but achieving widespread use in 1970 with the *System/370*.
(3) **A new partitioned architecture (for safer sharing of resources) and a new technology base**, including the introduction of LPARs and *PR/SM* in 1990.
(4) **A new hardware technology with a significant increase in core counts**, beginning with the adoption of CMOS-based[2] processor technologies in 1994.
(5) **The integration of open systems technologies** in 1999, in particular *Linux*, within established mainframe operating environments, thus creating a mixed-mode mainframe.
(6) **The stretching of the breadth of the mainframe with a hybrid[3] offering that includes tightly-coupled and managed Power Systems and Systems x blades**, with the introduction in 2010 of the *zEnterprise BladeCenter Extension (zBX)*. This was a fundamental redefinition of the scope of traditional mainframe computing.
(7) **The contemporary selling of zEnterprise as a cloud solution.**

Each of these seven cloud-related milestones from IBM's mainframe evolution to the present day will be described and detailed.

1. Setting the Stage – The Introduction of "General Purpose" Systems

In the early years, the "historical" mainframe values listed on page 1 initially were unrecognized because they were overshadowed by the dramatic scope of the System/360 announcement in 1964 and also because they could not be demonstrated or proven without an established track record (which would not come for several years). After all, the first S360 wasn't delivered to a customer until late 1965, a *Model 40*, I believe. Sales and marketing coaches told us to "sell the sizzle and not the steak" and there was plenty of sizzle to be had with IBM's announcement. First was the 50:1 performance range of the family, initially consisting of six models. Second was a full range of compatible and interchangeable DASD[4], tape, printers, and unit record I/O accessories[5]. Third, the systems were supported by common operating systems, compilers, and utilities, and any program written for any model was executable on all

[2] Complimentary Metal-Oxide Semiconductor.

[3] The use of the term "hybrid" in the context of cloud computing implies the blending of private and public cloud facilities; as will be explained later, I am using the term to apply to a uniquely mainframe context.

[4] Direct Access Storage Devices, as they were called for many decades. These were the "hard drives" in today's parlance, although some were actually "drums", with a head per track. Initially, these were physically huge and exceedingly costly per megabyte of storage (as a gigabyte was unthinkably large back in 1964).

[5] Like a card reader/punch, which generically were called unit record devices. For more background reading, see http://en.wikipedia.org/wiki/IBM_System/360#Unit_record_devices.

Exhibit 1 —IBM System/360 Model 40 Mainframe (circa later 1960s)

Above, the S/360 central processing unit is on the left against the wall, with the typewriter-like printer in the foreground on the left. Five tape drives are on the right in the rear, with two disk storage units in the foreground on the right; in between these is the card punch/reader. Collectively, when coupled with an operating system, these constituted the mainframe environment, although in today's parlance, the word mainframe tends to refer only to the central processing unit.
Source: IBM

others without modification. That still holds true today, any program written in conformance to the 1964 architecture will execute with no alterations required on any future mainframe, including today's *zEnterprise* system, the most current mainframe implementation.[6]

These were characteristics and virtues that had not been demonstrated by any prior computer system, whether from IBM or any others in the industry. The magic here was that the design clearly separated the implementation from the architecture. This is the architecture that prevails today and continues to protect IBM's mainframe customers' investments, both in hardware assets and customized software, year after year. Under the covers, each model was implemented quite differently, but all were completely backward compatible. At the low end, the *Model 30* had all of its instructions implemented in read-only-memory (ROM), i.e., microcode, essentially turning itself almost entirely into a channel during any I/O operation. The *Model 75,* which I first encountered at NASA Goddard in Maryland, was described to me as having every instruction "hard-wired". The intermediate models had progressive degrees of ROM and hard-wired instruction mix, the higher in the model range, the less dependent on microcode.

My Mainframe Indoctrination Began in 1965

My fascination with computer technology began in graduate school when I enrolled in a FORTRAN programming course in 1965 and was introduced to an *IBM 1620*. This was a small, relatively inexpensive computer introduced in the late 1950s that was aimed at the scientific and technical computing community, so by this time it was considered ancient. Many a student had run their first FORTRAN program on this machine, though during its lifetime of less than ten years its use went well beyond that of an instructional platform. It was simple, easy to use, took punched card input, and used a typewriter (pre-*IBM Selectric*[7]) output, and, even for its time, was notoriously slow with core memory cycle times in the tens of microseconds. That is three orders of magnitude slower than the memory technology in common use today. I was able to successfully code a regressive Simplex algorithm in FORTRAN II for a game

---

[6] Another significant factor was the ability to emulate the prior generation systems, i.e., the *IBM 1400* series on the *Model 30*, and the *IBM 7094* on the *Model 65*, thereby facilitating migration from these platforms to the S/360.

[7] For those readers too young to even recognize this term, the IBM Selectric typewriter was a technological innovation and marketing success from IBM. It had a golf-ball-sized print head that tilted and rotated replacing the standard typewriter's one striking-element-per-character/key design, which at higher typing speeds would often become entangled.

theory model that, thank goodness, required only one pass for a stable solution. I wanted more of this, much more.

It was at this time, having completed a stint in the military service and the completion of a Master's degree, that I was hired by IBM as a Systems Engineer (SE) and assigned to a branch office in the suburbs of Washington, D.C. This was October 1966 and IBM was recruiting frantically to staff the expansion of Data Processing Division (DPD) branch offices to meet the demand for the System/360. New managers were being promoted with only three or four years' experience in the business. These circumstances might suggest that the company had to lower their recruiting standards. I don't believe so based on the many bright and hard-working people I knew and worked with during this time; I felt very fortunate to be among them. Typical of this era, there were very few women in the branch offices, aside from administrators and secretaries, and those were limited to technical support assignments, such as mine. Female salespersons were still off in IBM's future.

On the first day, I was coached on the "dress code" (which officially did not exist) by a colleague. No big deal for me, as I had spent the majority of my last four years wearing a white shirt and a navy blue suit. On the second day, I was introduced to my first System/360, a Model 40 at a local university, similar to what is shown in Exhibit 1 on page 3. It would be an understatement to say that I was amazed beyond any expectation: the flashing lights, the disks whirring, the tape reels spinning with tape loops flying up and down in their vacuum columns, all while printers chattered away. These were new sounds to my ears and an example of the new pinnacles of technology that I actually could touch. It was a little overwhelming at first but I was hooked!

Six weeks later, I completed Basic Computer Training (BCT), which only meant that I was competent enough in 360 Assembler, FORTRAN, and COBOL to have my lab projects run to satisfactory completion. I also and had my first taste of the *Disk Operating System, DOS/360*. (The school also taught me to keep my wing-tips shined and my white shirt starched.)

My first assignment was to assist a more experienced SE with the upcoming installation of a Model 40 at university-affiliated research organization operating under the *Operating System/360 (OS/360)*, affectionately known as "Big Oz". My primary focus was to learn the intricacy of the operating system, which was just beginning to be understood by the field and, frankly, was a bit buggy, was quite slow on a Model 40, and had a voracious appetite for memory (measured in kilobytes, which is unfathomable currently). One reality of the times was that our customers' technical resources (systems and programming staff) were understaffed and were just beginning to get the appropriate training, mostly on the job working side-by-side with IBM SEs. In our lighter moments, we would sometime reflect that we were only one chapter ahead of them in the technical documentation. Nevertheless, as partners, we got the job done.

These were difficult times, not only for us but our families as well, with many long and late hours for all those working in the field, including the SEs that tended to the software and customer training and the

Exhibit 2 — IBM and the Seven Dwarfs

IBM was thought to be so dominant in the era following the announcement of System/360, that a now unknown industry observer, perhaps out of envy, coined the term "IBM and the Seven Dwarfs". The alleged seven were Burroughs, UNIVAC (Sperry Rand), National Cash Register (NCR), Control Data (CDC), Honeywell, RCA, and General Electric.

I remember well that the business was not so easily won nor to be taken for granted as some might assume. In my branch alone, I can recall a number of competitive losses to UNIVAC, CDC, and RCA, all from what had formerly been loyal IBM customers. When RCA and GE eventually left the field to concentrate on their core businesses, the remaining mainframe competitors were known as the "BUNCH", this being an acronym of the first letter of the names of the remaining five players. The Seven Dwarfs largely have vanished from the mainframe space, but IBM continued to face very formidable mainframe competition with which I personally accumulated a great deal of experience. Beginning in the mid-1970s, mainframes from Hitachi and Amdahl entered the market in the U.S. and Europe, joined by NEC and Fujitsu in Japan. Commonly referred to as the PCMs (Plug Compatible Mainframes), they were essentially software-compatible with IBM's systems and enjoyed great success for over two decades. By the mid-2000s, non-IBM mainframes all but disappeared in the U.S. and Europe.

Source: Clipper analysts

Field Engineers (FE) that were responsible for maintaining and keeping the hardware running, and also was true for the IBM laboratories and plants responsible for both. The IBM Corporation had bet more than $5 billion dollars initially[8], and much, much more through the years, and we were responsible for delivering on the promise that the company made to its current and potential customers The challenges and the rewards, both spiritual and material, were many. It was a very good time and a good place to be for a young man with a growing family; the future looked very bright, and this turned out to be true.

2. Large-Scale Virtualization Emerges

When announced in 1964, the S/360 could not be said to meet the definition of cloud computing on any dimension. But the vision was there, although certainly not articulated in cloud terms. After the announcement of the S/360, there was clear momentum building in the market, particularly among several leading universities, for a *time-sharing system*, ostensibly to support computing services to multiple simultaneous users. This culminated in the shipment in May 1966 of *the S/360 Model 67* and the *Time Sharing System (TSS/360)* monitor. A radically new design component, the *Dynamic Address Translation (DAT)* box was introduced, which enabled *virtual memory addressing*. This and other architectural extensions unique to the Model 67 provided users with a large address space, isolation of the operating system from the user's space, and most importantly, multiple address spaces, characteristics of what we expect from cloud computing today. The Model 67 was a giant step toward the ideals of cloud computing, time-sharing being the cloud computing paradigm of that era. But the system had limited success in the market; the design and architecture were years ahead of what the hardware and software technology could deliver and the system was withdrawn in 1977.[9]

With the introduction of *Operating System/Virtual Storage 2* (*with multiple virtual address spaces) (OS/VS2 (MVS))* in 1974, virtualization technology was expanded in support of the growing family of the *System/370 Advanced Function* systems.[10] MVS, which lingers on for some even today as the generic name for IBM's high-end mainframe operating system, introduced support of multiple virtual address spaces and the virtual memory facilities, as well as multiprocessing capabilities. Virtual addressing reduces memory fragmentation (therefore enabling a more efficient use of that resource), better isolation of tasks (jobs), provides a shared area for inter-job communications and, as a consequence, improves the system's internal security.

Also at this time, MVS introduced the prioritization of workloads, wherein users could define performance targets for programs, whether submitted as a batch job or one that was remotely entered. The degree of shared memory management and automation were modest but they were an essential step in the development of resource sharing and elastic capabilities needed for cloud-like infrastructure.

Win Some and Lose Some – It's What You Do After … That Counts

After completing a little over five years in technical support assignments, I came to the conclusion that future advancement in IBM with my skill set (not a programmer or engineer) would best be served with an assignment in sales. So in 1971, I joined the IBM team that was responsible for the Washington, D.C., municipal government. Through our influence, the powers-that-be adopted a plan that would unify most of their IT departments under the leadership of the Metropolitan Police and the DMV, though each department maintained its operational independence. This resulted in the elimination of a huge amount of duplication, established a basis for sharing data, leveraging of their best technical resources, and expansion of their services, all within the limits of the congressionally-mandated budget. Being a part of the planning and execution of this city-wide project allowed me to do some very rewarding pioneering work in database management, online transaction processing, and financial management systems. This provided me with the opportunity to lead my own territory.

Akin in some respects to today's public clouds, and a business model (rented capacity) that was

[8] For a retrospective on the conception and launch of the S/360, 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.

[9] The marginal success of TSS/360 opened opportunities for others, notably University of Michigan's *MTS*, Stanford's *ORVYL*, *WYLBUR*, and *MILTEN*, and McGill's *RAX* and MUSIC, which also were distributed by IBM. See http://en.wikipedia.org/wiki/Time_Sharing_Option.

[10] The System/370 featured dual-processor capability, main memory based on integrated circuits, and virtual memory addressing. The Models 135, 145, 155, 165 were the first out the door but were later superseded by the more powerful and flexible Models 138, 148, 158, and 168.

established many years earlier by IBM's Service Bureau Corporation (SBC)[11] and others, the Washington, D.C. area was fruitful ground for several companies that were created and sustained to satisfy the explosive needs of an expanding Federal bureaucracy. (During my tenure there, the Environmental Protection Agency (EPA) and the Department of Energy (DoE) were created.) One of my customers was a successful "outsourcer" for a federal agency, winning a contract in 1973 to provide, support, and manage a new data center to be equipped with two complete *S/370 Model 168* systems and a full array of storage and networking components. This was the customer's as well as my first experience with IBM's largest commercial system that required an external source of chilled water for its cooling. One issue that arose was ensuring that the installation of the chilled water system conformed to IBM's interface specifications. As a consequence of labor union restrictions, there was the question of who was to perform the actual physical connection of the building's chilled water lines to the system's cooling unit's intake manifold. IBM's field engineers were loath to allow a plumber, though likely a skilled mechanic, access to the inside of one of their machines. We worked out a compromise: the plumber would handle the supply lines, IBM's FEs would handle the manifold's connectors, and the drinks were on me.

A few years later my career in IBM received a true test of its resiliency with a critical situation at another of my customers. This company was a successful a supplier of advanced time-sharing services and application primarily to the financial and manufacturing sector. Several of its managers had IBM roots, but this did not assuage them from considering the installation of an *Amdahl 470V6*.[12] Although its services were hosted and backed up by a pair of *System/370 Model 155*s, their operating system was based on an early and highly modified version of OS/360. It also was decided at this time to eliminate the separate backup system, as the customer was confident of the Amdahl system's reliability and had more than enough scale to meet their foreseeable needs. Moreover, the company was not dependent on any IBM-provided software. Despite all our efforts (with lots of help from sales management and Poughkeepsie-based engineers[13]), IBM lost the business. The fallout on this deal was huge, personally and professionally, and, as could be expected, it shook my confidence. The new Amdahl, S/N 00012, was the first of its kind in our region, and soon proved to be a stepping stone for the Amdahl Corporation eventually winning a significant share of business not only with the Federal government but in many other sectors as well.

I wish I could report that this customer had made a poor business decision, but I can't. The Amdahl 470V6, with its bright red panels, served this customer well for several years to come. However, one thing the customer miscalculated was the impact of having only a single system on their development operations. Their programmers had become accustomed to running development and test whenever they wished. However, for reasons of reliability and performance, the programmers now were confined to the second and third shift and even that was sometimes restricted. Their advanced development and sometimes even routine maintenance was stifled. This became an opportunity for me to propose an *S/370 Model 148*, which would more than meet their technical needs, was affordable, and could perform as a backup system, albeit limited in its scope. Their systems' nicknames became "Big Red" and "Little Blue". The regional manager that had "ripped" me badly for the competitive loss in the prior year was happy, and with some humility on his part, awarded me generous recognition a year later. One of the key factors to our winning strategy was introducing the company to the notion of adopting ***VM/370* and its virtual machine architecture as a fundamental element of its software operations**. This was a difficult task because the technical management and staff were invested heavily and emotionally attached to their current unique, powerful, and very flexible, operating environment. They were not interested, nor were they inclined to trust, software that emanated from an IBM laboratory. But VM/370 was to test all their assumptions and would win their hearts.

VM – the Roots of the Mainframe Cloud Take Hold

One of the key factors to IBM's progressively innovative strategy was the introduction of VM/370 and its virtual machine architecture as a fundamental element of its mainframe evolution as a cloud

[11] In early 1973, IBM sold its SBC operations to CDC for terms very favorable to the latter as part of the settlement of an anti-trust lawsuit that alleged IBM had introduced a "phantom" supercomputer (for those times) that it never planned to build in order to forestall orders for CDC's products. See Notes at http://en.wikipedia.org/wiki/Control_Data_Corporation.

[12] The Amdahl Corporation was founded in 1970 by Gene Amdahl, the chief architect of the System/360, who allegedly left over a dispute over the direction of IBM's large systems architecture. In 1975, the company announced its first system, the 470V6, in response to IBM's S/370 virtual memory systems. It more than matched the performance of the IBM S/370 Model 168 and was offered at a significantly lower price. See also Exhibit 2 on page 3 above.

[13] I am still acquainted with two of these gentlemen.

platform. Virtualization of system resources is key to their most efficient use and contributes greatly to the flexibility and resiliency of today's cloud infrastructure. As was stated earlier, when S/360 initially was announced, there were no provisions to accommodate a robust time-sharing environment. IBM's Cambridge Scientific Center in Massachusetts took on an effort in 1966 with several goals,[14] including the development of an efficient time-sharing system for *internal* use.

Using an address translation mechanism designed by Gerrit Blaauw, one of S/360's principal engineers, for a Model 40 as their first research platform, two independent, and sometimes confused, software components were developed.

- ***CP-40 – the Virtual Machine Control Program*** – a virtualization hypervisor that is able to create a virtual machine (CPU, memory, mini-disks etc.) that replicates the S/360 architecture. A copy of OS/360 or DOS/360 could run on these virtual machines quite transparently.[15]
- ***CMS – the Cambridge Monitor System*** –which ran in its own CP machine and served as the the interactive component and the user interface. Initially, it supported development using Assembler, PL/I, and FORTRAN whose programs could run on CMS or OS/360, a line-mode editor, text processing, a command processor, and utilities to manage the environment.

A version was created for the S/360 Model 67, called *CP-67/CMS*, was released with limited support in the spring 1968, and was offered commercially by several companies as a time-sharing (fractional use) basis. Most importantly, it became very popular throughout IBM internal sites. It was IBM's first "open" system, distributed freely in source code format, and it was user-friendly even before decent green-screen terminals were generally available. I had my first personal experience with this system in early 1969 when I attended an IBM time-sharing systems class in Poughkeepsie, later to become my home town. We were given the opportunity to learn the fundamentals and get some hands-on with *TSS/360* and CP/CMS. Even on a dial-up connection to the Cambridge Lab, CP/CMS impressed us; it was responsive, easy to use, and it never (well, rarely) crashed. We wondered why IBM was keeping it buried in the lab. Instinctively, we knew that our customers would love it; no market research was required. Finally, when IBM announced virtual storage for S/370 in August, 1972, *VM/370* joined *OS/VS* and *DOS/VS* in support of the S/370 family. VM's future seemed secure.

There was a period that followed in which the forces that controlled the roles and destinies of MVS and VM, respectively, appeared to be in conflict. My personal perspective is that there were several issues, starting first with level of investment in development, the technical scope of each system, and the appropriate role for each in customers' mainframe data centers, which included the aspect of what each would contribute to the sales of new or increased investment in mainframe hardware. The MVS people seemed to have the stronger position, if only judged by the level of investment and the size of the program and its extensions already produced.[16] Furthermore, customers were becoming confused over the roles each would play in their data center and how it would affect their resources. At the same time, pressure also was coming from internal sites for continued investment in VM particularly because the role it played in in supporting thousands of IBM users throughout the world who depended on it. For instance, IBM's first e-mail and calendaring systems was based on *VM/PROFS* (Professional Office System), which later became very popular with a number of customers. We know now that these issues were resolved to the general satisfaction of all.

VM has continued to grow and support the mainframe as it has evolved over five decades to its current state, from VM/370 in the 1970s, *VM/SP* and *VM/XA* in the 1980s, *VM/ESA* in the 1990s, and now *z/VM*. *Why is this important for us to understand?* An essential premise of this paper is the evolution to a mainframe-based cloud environment would not have been possible without VM. VM, as it evolved through several architectural extensions, provided for a significantly higher degree of flexibility to the mainframe, enablement of self-provisioning, responsiveness to individual user demands, and more effective resource utilization of the mainframe, thus extending it into environments and application arenas never foreseen by the original designers.

Today, you can't imagine operating a shared cloud without it being virtualized; that's the only way to

---

[14] The primary source for many of the details in this section are extracted from a presentation coauthored by Erich Amrehn and James Elliott, which is available at https://www-950.ibm.com/events/wwe/grp/grp019.nsf/vLookupPDFs/7 - VM-45-JahreHistory-EA-J-Elliott [Kompatibilitätsmodus]/$file/7 - VM-45-JahreHistory-EA-J-Elliott [Kompatibilitätsmodus].pdf

[15] This is akin to running Linux in a VMware partition.

[16] There was also talk of the reverse scenario – one in which the MVS would become a subset, a dependent of VM. In many of today's modern mainframe environments, it would appear that a synthesis of both worlds has been achieved quite amicably.

do it. The long-standing champion and enterprise standard for all things virtual is the IBM mainframe. At this time , the "cloud" moniker had not yet been invented, but the origins of today's virtualized world can be traced back to the 1960s and to many concepts (like time sharing) and features (like virtualization) added to the later members and successors of the System/360.

3. A New Partitioned Architectural and Base

In early 1979[17], I moved on to the IBM Poughkeepsie Laboratory to become closer the development and market planning of mainframe technologies.[18] This was a time of great transition once again. The System/370 family was getting long-in-the-tooth even after having been extended by the introduction of the *303X* family[19] in 1977. Though they were largely based on the S/370 Models 158 and 168, a number of enhancements resulted in significant performance improvement, larger memory, and integrated channels that dramatically reduced floor space and other environmental factors. It was largely a stop-gap program because the follow-on program, known as *Alpine* then *Adirondack*, was significantly delayed; another proof-point that scheduling invention was very risky and often led to disappointment. Finally, in late 1980, the *S/370 3081 Processor Complex* was introduced with a great deal of fanfare. The "sizzle" was embodied in the TCM (Thermal-conduction Module)[20], a dyadic (dual) processor, a revolution in architecture and design, including more extensive use of microcode, circuit integration, packaging, cooling, and manufacturing technologies that resulted in vastly improved performance (up to 21 times that of 3033), a compact footprint, greatly improved energy and cooling efficiency, reliability, and serviceability. This family was extended over the next few years to include several models of the *3081*, the uni-processor *3083* models, and eventually the *3084*, a quad processor.[21] In 1984, the S/370 architecture, the high-end 308X and the midrange 4381 midrange processor families, all were enhanced by extending application addressability from 24-bits to 31-bit, and thus from 16MB to 2GB. This was supported by a new version of the operating system, *MVS/XA*.

The year 1990 marked another major architectural extension, *LPARs* (Logical Partitioning) and *PR/SM* (Processor Resource/Systems Manager) and also with the announcement of the *System/390* and the *ES/9000* processor family[22], the high-end supported by a new version of the operating system, *MVS/ESA SP (MVS/Extended System Architecture-System Product).* Each LPAR operates like an independent system running its own operating environment and is essential to the management of ever-increasing processor capacities, thus an essential element of the cloud computing paradigm. The PR/SM function is implemented in the hardware as a Type 1 hypervisor and is the means of enabling LPARs to share and dynamically add or delete physical resources (including processors, memory, and I/O).

Together, LPARs and PR/SM created the virtual environment that is familiar today – the one that other systems want to be when they describe their server environments as "mainframe-like". In reality, none do it as well as the IBM mainframe at the beginning of the 21st century.

4. New Hardware Technologies Increase Core Counts

In 1992, I was working with the engineering team to perform a market assessment for the next big S/390; we called it *H8* internally, the eighth iteration of the high-end. The early H6 systems already were in the field, ranging from rack-mounted entry *Model 120* through the *Model 900* six-way MP behemoth, all with upgrades in the development pipeline. The engineering team was faced with several problems for the next generation, problems that were well-known and characteristic of the Bipolar TCM

[17] Even though I had moved out of sales and into development, I was still able to attend my sixth One Hundred Percent Club with my former teammates from Washington, D.C., in late April in New Orleans. The event was marked by spilling a bowl full of spicy-sauced shrimp on my suit pants. Fortunately, I had a back-up.

[18] Part of my early education was to discover how well the Poughkeepsie Lab kept track of its largest mainframe customers and their installed capacity. Thus there came to be the institution (unofficial, of course) of the "Ten-MIP Club", customers with a total capacity of at least 10 MIPS in their data center. At this time, there were less than 12 customers with that distinction. Oh, have times changed!

[19] They were the 3031, 3032, and 3033, in order of increasing performance. The last system I was to install as a salesman was a 3032 – S/N 00002.

[20] It included innovative use of multilevel ceramic substrates, helium, bees wax, and spring-loaded pistons capped by a water jacket.

[21] I worked on the business case for the 3084Q, which was not justified by its initial forecast (only 5) but by broader customer demands for a larger scale system. IBM sold more than 50 over the product life.

[22] Initially, the family included entry level rack-mounted, midrange air-cooled and high-end water-cooled systems up to 6-way multiprocessors. Optionally, vector processors were available which enabled supercomputer-level (for that time) scientific and engineering applications.

(Thermo-Cooled Module); it was fast, but had ever-increasing power requirements and thus presented more complex cooling challenges. On top of these factors, the circuit count was extremely high and was a severe challenge to testing capabilities; it was believed that producing a module that met IBM's increasingly stringent reliability criteria for its mainframes might not be attainable. Furthermore, the competitive assessment of Hitachi and Fujitsu PCMs (see Exhibit 2 on page 4) indicated that they were capable of delivering their next generation products with performance that could exceed what IBM had in its committed plan.

IBM now was forced to make a decision that I believe, in many respects, rivaled the decision in the early 1960s to invest billions of dollars in the System/360. IBM had to move to a new technology base, Complimentary Metal-Oxide Semiconductor (CMOS), which substantially would address not only the cooling and power consumption issues, but escalating cost issues as well. CMOS technology had been around since the early 1960s and was well-understood by IBM's engineers, but until the 1990s had not been seriously considered for implementing a high-performance mainframe. There remained one very large, very important issue that made the decision so risky for IBM: it would take at least four years of development before a CMOS-based multiprocessor system could reach the single-system capacity (about 1000 MIPS) of the current bipolar-based systems. In April 1994, IBM announced the first generation of the new family, *System/390 Parallel Enterprise Server* models. The scope of the announcement included the *S/390 Parallel Transaction Server (HPTS)*, the *S/390 Parallel Query Server (HPQS)*[23], a high-speed coupling facility, a new version of MVS/ESA, and a number of other software enhancements, including DB2. This was a very significant architectural extension that was designed to enable up to 32 system images in the Sysplex and by doing so dramatically increased the scalability and availability characteristics of the platform.[24] But it would not be until 1998 with the delivery of the G5 models, with up to a 10-way (10-core) multiprocessor, that IBM was able to deliver at the performance level of its prior family and match that of its mainframe competitors. In the ensuing time, IBM's high-end mainframe market share tumbled; every customer's mainframe acquisition was being contested, exacerbated by the fact that many clients were finding it advantageous to adopt a multi-vendor strategy (having at least two brands of mainframe in a data center) – all the better to keep each vendor on their toes with their pencils sharpened for the next deal.

The adoption CMOS-based processors as the building block for the midrange and high-end mainframe systems and continued architectural enhancements had enabled IBM to drive down the cost of computing and dramatically increase its scale by not only addressing manufacturing costs but also cost of the data center infrastructure, in terms of floor space, power, and cooling, required to accommodate it.[25] IBM also had designed its systems to be easily upgradable by the use of common building blocks, pluggable (attachable) components, and microcode-enabling of additional capacity and new functions, thus eliminating the need to disrupt normal operations. It was during this period of time that the industry-leading *IBM S/390 CMOS Cryptographic Coprocessor* chip was introduced. It received the U.S. government's highest certification for commercial security at that time, a FIPS Level 4 certification.

Taking a Breath to Catch Up With Intervening Career Changes

It was at about this time that our department switched its focus from strategic market planning toward the establishment of a market research department. The impetus was provided by a sister division embarking on a project to build a market segmentation model for their products. The Data Processing Division (or whatever we were called at that time) had to have one, too. One of the Big 5 market consultancies was engaged to direct us in building a model that, ostensibly, would assist the division in the development and marketing of more targeted products. After an effort of close to a year, we ended up

---

[23] HPTS and HPQS were focused offerings in response to *Tandem Computers* (see http://en.wikipedia.org/wiki/Tandem_Computers and *Teradata Corporation* (see http://en.wikipedia.org/wiki/Teradata) offerings, respectively. I worked as a marketing team leader for HPQS for almost a year. About five customers took delivery, but found it could not meet their performance objectives. Within a year or so systems were converted to standard configurations with the assistance of IBM software and hardware engineers. Within two years the designation, HPTS and HPQS were abandoned to focus on general purpose systems.

[24] New invention was again required in both software and hardware, not the least of which was the coupling facility that synchronously linked processors and allowed them to share system management and logging data. The final design was based on a modified version of a 43XX midrange processor; however the current implementation uses a single core of the System z 6-core chip, though customarily duplexed or triplexed for availability. The capabilities of the Parallel Sysplex facility continue to be enhanced and form the basis of IBM's mainframe real-time backup and recovery systems offerings.

[25] For example, the need for an external cooling unit connected to an external source was eliminated. Thus, it was not uncommon for customers to have the benefit of delaying or completely avoid the expansion of their data center facilities even while growing their capacities at double-digit annual rates.

with four or five binders of contractor-prepared presentation material and a pile of analytical listings. The effort was instrumental in building awareness of market segmentation models and how segments could be differentiated but was of little use for guiding our strategies, because of its complexity. The people who built it understood it; the people who had to use it did not.

This did not dim our hopes for producing useful market segmentation models, but with more limited goals. I became the project manager of building preference models that would be capable of projecting market share and potential revenues. The first was focused on the high-end mainframe choices. The model met its goals and reinforced our confidence in for preference models. The second preference model that I managed made the startling revolutionary concept of including non-mainframe customers in the surveyed population. Its primary contribution to our understanding was that it focused on the fact that the mainframe was and would continue to share the market with distributed systems; client-server was the *au courant* terminology at the time. I spent about four years working with these models and received a Division Manager's Award for my efforts. As a bonus, I also got much better acquainted with Silicon Valley, the home of our contractor. Meanwhile, the market research department was working on a number of other projects, the most significant of which was to build a global longitudinal model (that is one the tracks over several time periods) that would test a number of aspects of customers' awareness, attitudes, intent, and behaviors with respect to the mainframe divisions' products and strategies. Refreshed semi-annually, it provided a wealth of information and trend data that became essential to the division's executives, developers, market planners, and sales teams. For that effort, our department and each of us individually received a Corporate Market Research Award for outstanding contributions to the business. This research effort continues to this day.

5. The Integration of Open Systems Technologies Gives Birth to the Mixed-Mode Mainframe

While a student at the University of Helsinki in 1991, Linus Torvalds began development of an operating system kernel as a POSIX-compliant[26] free alternative to the commercial distributions of UNIX which were subject to licensing fees. Known as *Linux*, an amalgam of his first name and UNIX, it was originally developed for Intel x86-based computers, but quickly spread to other platform architectures. IBM's lab in Boeblingen, Germany, began development on a collection of patches and additions to the current Linux kernel to port it to the S/390 architecture and was released to interested parties in December 1999. This was quickly followed by the development and release of the *Integrated Facility for Linux (IFL)*, to be available as special feature of the then current G5 and G6 high-end mainframe systems. The IFL processor was limited to run only under Linux or VM/ESA, but with no degradation of speed, and was priced well below a standard (unrestricted) S/390 processor (which otherwise was the same).

By a stroke of genius, or just a set of fortunate circumstances, IBM's mainframe engineers had created a mixed-mode platform, i.e., a system that could run multiple architectures simultaneously, transparently, and interchangeably.[27] With Linux running on a mainframe core, a vast new universe was opened up for the mainframe and allowed it to participate robustly in the open systems world dominated by x86 architectures and open systems programming models. An IFL running VM enabling multiple *Linux on System z* virtual machines (images) easily can be recognized as a cloud in its own right, with each image simulating the whole system. VM with Linux for System z took the best of the mainframe and made it available to all Linux applications and users. As the acceptance of the Linux for System z became evident, IBM announced that it would be investing $1B in Linux development over the next several years, including a well-staffed organization to support its extensions for the mainframe and its other system platforms.[28] Clearly, enhanced cloud-like thinking was present in many ways on the IBM mainframe.[29]

[26] **P**ortable **O**perating **S**ystem **I**nterface", a family of standards specified by the IEEE for maintaining compatibility between operating systems. POSIX defines the application programming interface (API), along with command line shells and utility interfaces, for software compatibility with variants of UNIX and other operating systems. See http://en.wikipedia.org/wiki/POSIX.

[27] IBM mainframes had been running multiple operating environments simultaneously for some time, but these tended to be IBM-provided mainframe operating systems. With the addition of Linux, an open environment, mixed mode took on a wider and more important meaning. Some called this "hybrid processing", although I will reserve that term for a later milestone.

[28] As a member of the mainframe forecast team we had no idea that Linux for System z would become so popular as there were no precedents to use for guidance; it turned out that the optimists were correct, though I think the level of acceptance even surprised them.

[29] During this time, I was still in the Market Research Department and I took on responsibility for a project for the Linux marketing leadership. We tested a large population on their awareness, attitudes, and intent with respect to IBM's Linux offerings. The resultant data provided a firm foundation that confirmed the direction of the current Linux development and assisted in the guiding of

In October 2000, IBM announced the first of the *zSeries* mainframe systems, the *z900*, which notably extended application program addressability to 64-bits and 16 exabytes of address space, all while remaining upward compatibility with 24-bit and 31-bit programming. The z900 was capable of expansion to 16 processors[30], including any mix of standard z/architecture processors and IFLs. This was followed by the *z/800* in February 2002[31], a 4-way processor, including an IFL-only model that would run only z/VM and Linux. This action was a significant and clear demonstration of IBM's commitment to lowering the cost of computing, expanding into new application domains outside traditional mainframe applications, and broadening the appeal of the zSeries to non-mainframe clients, particularly those who wanted to run many Linux images. At a very attractive TCO, it offered an ideal platform for consolidation of Linux workloads particularly those running on x86 servers typical of small to medium cloud implementations. The need to provision and de-provision Linux images easily and quickly is at the heart of the requirements for a compute cloud. It now was possible without a lot of traditional mainframe knowledge or skills.[32]

Virtualization is "Real"

A mixture of amusement and dyspeptic agitation befalls me whenever I read something about the miracle of virtualized environments among the minions of the x86 devoted as the solution for everything. Every once in a while they need a reality check that IBM has been doing this on mainframes for almost 50 years. It is so imbedded in the mainframe experience that we take it for granted.

Virtualization is not optional on System z mainframe. The first layer, an imbedded hypervisor if you will, is provided by *PR/SM*, which is the means to deploy LPARs, up to 60 in the latest generation *zEnterprise EC12*. LPARs are defined with a mix of virtual (or dedicated real) resources: processors, memory, network, I/O devices, and a priority. Importantly, these resources may be shared dynamically and the total of the virtual resources allocated may very well exceed the total real resources of the system. This is central to essential nature of mainframes being capable of continuous running at full capacity whenever demanded regardless of application workload mix. It also is essential to the ability to recover very quickly, essentially transparently, if any of the system's resources become crippled or unavailable, whether planned or unplanned.

Any of System z's supporting operating systems can be loaded into an LPAR; one or more copies of *z/OS*, *z/VSE*, *z/TPF*, or z/VM. A double tier of virtualization is provided by z/VM when its virtual machines are defined as z/OS, z/VSE, z/TPF, or Linux for System z, the latter generally defined to execute only on the IFL processors[33]. Within the context of z/VM and Linux for System z, many hundreds of Linux server images may be deployed, limited only by the real resources available to that LPAR pool, there being no architectural limit imposed either the System z hardware or software.

Evolution of the Mainframe Family Continues

Following the zSeries z900 and z800, the *z990* (a.k.a. "T-Rex"[34]) and *z890* were introduced in the springs of 2003 and 2004, respectively. The high-end z990 expanded the multiprocessing to 32-way (the z890 was still limited to 4), and up to 256GB of main memory. A second specialty processor was added to the mix, the *zSeries Application Assist Processor (zAAP)*, which was designed to accelerate JVM and some XML System Services functions, consequently off-loading such workloads from running on standard processors and z/OS, thereby lowering the total cost of ownership. In 2005 and 2006, IBM introduced the *zSeries z9 EC (Enterprise Class)* and the *z9BC (Business Class)*, respectively. Building on the

---

future strategies. I had the opportunity to attend a Linux developers conference in NYC, met a number of very interesting folks, and visited a number of vendors' booths collecting an array of tchotchkes in the process. And I could hardly spell Linux.

30 Over the life of the mainframe the use of the words "processors", "engines", and the more contemporary, "cores" have been used somewhat interchangeably. They all refer to the most discrete physical computational element within a computer.

31 The Clipper Group has published numerous papers on IBM systems, storage, and related technologies. The first of the zSeries is cited here. See **The Clipper Group Navigator** entitled *IBM Takes The Wrap Off of Raptor – Smaller zSeries Mainframe Targets New Users*, dated February 18, 2002, and available at http://www.clipper.com/bulletins/2002/z800_final.pdf.

32 By this time, I transitioned into a new department and became the division software forecaster. My first project was to develop the forecast for *z/OS.e.* The purpose of this version of z/OS was to offer very much lower entry price for the operating system, particularly to support the z/800. In order to do this, it was limited for use of "new workloads" only, therefore it would not execute programs written in COBOL, PL/I, or FORTRAN, or host CICS or IMS applications. In my tenure as a forecaster, I shared the responsibility for the development of statistically-base forecast models for both mainframe hardware and software and pricing models for new offerings.

33 This is an economic choice as IFLs are priced considerably lower than standard z/architecture processors as stated earlier.

34 A dig to those who for years prior had been referring to the mainframe as a dinosaur – big and dead.

prior generation, upper capacity was expanded to 54-way, granularity options were expanded in the lower capacity models, plus there were a number of improvements in workload management and dynamic resource allocation features that greatly improved the cost-effectiveness of the platform. During this period, another specialty processor, the *zSeries Integrated Information Processor (zIIP),* which was designed to offload, in the same way and effect of the zAAP, certain parts *of DB2 for z/OS* DRDA, star schema, and IPSec (Internet Protocol Security) workloads.[35]

In 2008, IBM produced the double whammy by introducing the *zSeries z10EC* in February followed by the *z10BC* in October. The z10EC processor speed was improved to 4.4GHz, ten additional user configurable processors were added scaling in five models from 12-way to a maximum of 64-way, resulting in over 30,000 MIPS total capacity. Main memory range was extended to 1.5 TB. There were on-chip processors for data compression, cryptographic, and decimal floating-point operations and 50 new instructions were added, mostly focused on compiled code efficiency. Also introduced was *HiperDispatch*, which all but eliminated data latency between LPARs. In a consolidation model created by IBM comparing the resource requirements for a 3200 virtual machine workload running Linux on Intel servers, the z10EC was 15 times more space efficient and over 20 times more power efficient.[36] The Business Class system introduced later in the year was designed for simplicity, flexibility, and lower cost. Up to five 3.5 GHz processors could be configured as standard z/architecture processors yielding over 2760 MIPS, but could also be configured in a range of sub-capacities as low as 26 MIPS, a fraction of one processor. Five additional user-configurable processors were available for specialty engines, including IFLs, zAAPs, and zIIPs. Also, a raised-floor was not required for installation. Contemporaneous IBM-sponsored surveys indicated that System z was actively supporting web serving, web application serving, data serving, and development at more than half the installations surveyed. Business intelligence, e-commerce, and network serving applications also were common. These workloads are typical of those that ordinarily might found in distributed (x86) server implementations of a public or private cloud.

The offering of a mixed-mode mainframe platform capable of running of Linux for System z in an LPAR while running native mainframe operating systems and, more importantly, the IFL offering with its use of z/VM as the cloud manager for many virtual Linux images running under its control are where the mainframe can be said to have become an open cloud computing platform. While this all depended on decades of mainframe inventiveness and underpinnings, to the less-informed outsider, the IFL probably is the clearest mainframe cloud milestone.

6. The Hybrid Mainframe Redefines Enterprise Computing

It would be fair to say that even the closest of mainframe watchers were caught by surprise when he *zEnterprise 196 (z196)* was revealed publicly in July 2010. At this point in my professional career, I had been retired from IBM for a little over a year and now was an industry analyst for The Clipper Group. However, while still with IBM in my last assignment in the department that was responsible for mainframe pricing, I was well aware of the products that were in the development stream. It was recognized that pricing a product that had its roots in mainframe technology but which also integrated non-mainframe technology was going to be challenging. The pricing models for both hardware and software were quite different in these domains, so the result had to make sense in both worlds and not become an impediment to customers' consideration. So I was pleasantly surprised when I attended the zEnterprise announcement event in New York City.[37] It was not to be the "z11" that many had expected. The consensus reaction was that only IBM had the vision, the technology, and incentives to take so innovative a step. The zEnterprise system reflected the acknowledgment of several key facts that had been emerging in the IT industry: First, modern data centers consist of, and will continue to be for the foreseeable future, multiple platform architectures. Second, clients have been struggling with understanding the means, technically

[35] In later generations more functions were included including incorporating designated zAAP workload onto the zIIP (zAAP on zIIP) in 2009. By using these specialty processors the performance of the targeted workload is significantly improved and the cost-effectiveness of the platform is enhanced because these processors are priced much lower than the standard processors, and z/OS usage, which is priced by capacity, is lowered.

[36] In this comparison, a z10EC with a 64-way IFL configuration is capable of running at modest utilization rates the same workload as an array of 1600 x86 cores running Linux. It is obvious the manpower requirements to manage each of these would vastly favor the mainframe. In addition, the z10EC would have an availability of 99.999% or better (the equivalent of 5 minutes downtime in a year running 24x7); this rate is highly unlikely for an x86 configuration of this size.

[37] Me and several hundred others.

and financial, to deal with the plethora of technologies and their integration. And third, the most rational and effective solutions lie in focusing the management and control of the required resources around their core systems. The announcement of the zEnterprise system, beginning with the high-end z196, was the IBM's first step in addressing these issues with its mainframe offerings. Existing cloud infrastructures on distributed systems, particularly those that required fast, reliable, and secure connections to mainframe processes, could be incorporated under a common end-to-end mainframe management envelope.

The *zEnterprise 196 (z196)* central processing unit[38], with a choice of five models, consists of up to ninety-six 5.2 GHz processors, the fastest in the industry to date, eighty of which were user-configurable and delivering a total usable capacity exceeding 50,000 MIPS. The processor chip included numerous cache enhancements, added 100 new instructions, a Super-scalar pipeline, and Out-of-Order (OOO) instruction provided substantial improvements to CPU-intensive, *Java*, and *C++* applications, along with traditional workloads including batch and transaction-oriented systems, such as *CICS*, *IMS*, and *DB2*. Moreover, the z196 delivered its potential in an environmental footprint similar to its predecessor, the z10EC.

In addition, the z196 could be coupled through two high-speed internal links, one for data and one for control, to a *zEnterprise BladeCenter Extension (zBX)*. The zBX could be populated with either *IBM POWER* blade servers or *IBM System x* (x86 architecture) blade servers or a mix of both up to a maximum of 896 cores. The POWER blades operated with *IBM AIX* while the System x blades operated with Linux. The third element of the IBM zEnterprise System was the *zEnterprise Unified Resource Manager*, or *zManager*; this was the glue, the software management piece that unified the processing resources of the z196 and the zBX through its hardware, platform, and service management functions. The zManager was distributed as an extension to and operates through the *zEnterprise Hardware Management Console (HMC)*. Specifically, the zManager provided:

- Hypervisor management and creation of virtual networks.
- Operational controls, service and support for hardware/firmware.
- Network management of private and secure data and support networks.
- Energy monitoring and management.
- Workload awareness and platform performance management.
- Virtualization management – a single view of virtualization across the platform.

Of course, all of these are important to running enterprise-class clouds of computing.

The attractiveness of the z196 was in the system being designed and built using well-known and well-understood but heterogeneous computing elements, where each brought its unique strengths to bear on the common goal, that being to provide the most optimized mix of power, efficiency, and resiliency to the delivery of information technology services. The need for both vertical and horizontal scale as well as the diverse needs for specialized services is well-fulfilled with this new IBM hybrid model. The combined power and virtualization capabilities of almost 1000 cores (across the z196 and zBX blade servers) under unified and secure management truly were revolutionary. **With this announcement, IBM brought forth and embraced its vision of an integrated mainframe-based hybrid system as a true fulfillment of the cloud computing model on the mainframe platform**.

In July 2011, IBM extended the zEnterprise hybrid model to the midrange with the announcement of *the zEnterprise 114 (z114)*[39]. The system included all the elements of the z196 announcement a year earlier, including the zBX and the zManager, but limited to a total standard capacity of 3100 MIPS. This was delivered in two models with up to 10 user-configurable processors plus the capacity and flexibility offered by the POWER or System x server blades within the zBX.[40] The year was capped with the announcement in October for the support of *Windows Server* on the zBX System x blade servers, in addition to the existing support for Linux. With this much anticipated enhancement, the "window" of opportunity to address flexible, resilient, and secure Cloud computing requirements expanded dramatically[41].

---

[38] More detail can be found in **The Clipper Group Navigator** entitled *The IBM zEnterprise System Reaches Out — Higher, Wider and Deeper*, dated July 22, 2010, and available at http://www.clipper.com/research/TCG2010033.pdf.

[39] For more detail on the z114, see **The Clipper Group Navigator** dated July 12, 2011, entitled *IBM zEnterprise in the Midmarket - Revolution or Evolution?*, which is available at http://www.clipper.com/research/TCG2011024.pdf.

[40] The Enterprise Linux Server, an all IFL configuration with supporting software, and pricing model was introduced. The ELS will be discussed further in a later section.

[41] For more details, see **The Clipper Group Navigator** entitled *IBM's zEnterprise Really Stretches Its Boundaries – New Windows Are Opened*, dated August 12, 2011, and available at http://www.clipper.com/research/TCG2011034.pdf.

Exhibit 3 —Current-Day IBM zEnterprise Mainframe (zEC12) with Tightly-Coupled zEnterprise BladeCenter Extension (zBX)

Source: IBM

The Current Generation Continues the Mainframe Transformation

In August 2012, IBM announced its latest and the current high-end mainframe, the *zEnterprise zEC12 (zEC12)*[42], adding significantly more capacity and incorporating a number of technical advances. A new 6-processor (core) chip running at an industry-fastest 5.5 gHz was introduced, becoming the building block for 120 active cores, 101 of which were user-configurable. Its total capacity exceeded 78,000 MIPS and was able to deliver this in essentially the same power/cooling/footprint as the z10EC, two generations earlier. Again, the instruction set was expanded, on this occasion to enhance Decimal Floating Point operations. A number of innovative features were introduced that resulted in significant improvements in the system's availability, security, and reduced data latency. (See Exhibit 3, above.)

In July 2013, IBM announced the twelfth generation midrange, the *zEnterprise BC12 (zBC12).*[43] Inherited from its big brother was the six-core chip running at about 75% of the rate on the zEC12. With a single-engine MIPS rate of 1064, a z/OS maximum configuration of six CPs delivered a total capacity of 4900 MIPS, expressed as z/OS capacity. With a maximum configuration of 13 cores, the actual maximum capacity of the zBC12 processor is a good deal larger, the equivalent of about 8700 MIPS, when the computing capacities of the seven special-purpose processors are included. With a midrange mainframe entry price of less than $100K, it had the potential to grow to a 13 IFL configuration hosting z/VM and Linux for System z. Included were support for Microsoft Windows Server 2012 (on zBX), enhanced workload policy-based performance management for AIX, Linux on System x, and Microsoft Windows, and improved availability across the zEnterprise ensemble with alerts and notifications for zBX virtual blade servers in remote recovery configurations.

Data Analytics Solutions Means Get Answers Quicker

Along with the zEC12 announcement in 2012, IBM announced enhanced capabilities for accelerated, real-time analysis of mainframe data without having to duplicate or move it from where it naturally resides in DB2 databases via the *IBM DB2 Analytics Accelerator for z/OS Version 3 (IDAA V3)*[44]. Powered by *Netezza* technology, the *Accelerator V3*, as it is known, is a workload-optimized appliance

---

[42] For more detail on the zEC12, see **The Clipper Group Navigator** entitled *The IBM zEnterprise EC12 - Bigger, Better, Faster*, dated August 28, 2012, and available at http://www.clipper.com/research/TCG2012019.pdf.

[43] For more detail on the zBC12, see **The Clipper Group Navigator** entitled *IBM's zEnterprise BC12 - More of What You Want*, dated July 23, 2013, and available at http://www.clipper.com/research/TCG2013013.pdf.

[44] The IDAA V2 was announced in October 2011 and was available for attachment to the z196 and z114. It provided essentially the same function as the later version.

specifically designed to assist enterprises meet the challenge of complex analytical needs. It can lower the turnaround time and predictability of complex queries dramatically; 1000 or more times improvements having been observed in real world applications. The Accelerator is directly attached through high-speed links to any zEnterprise processor as a logical extension of DB2; query tuning is eliminated as well as the need to create and maintain indices. The Accelerator is completely transparent to the user and, therefore no query modifications are required to begin to exploit its performance.[45]

Mainframe users commonly accumulate huge volumes of transaction data, usually associated with services about their clients and customers, and log data, usually associated with data base management and recovery systems. Gleaning intelligence from this unstructured data is an enormous task. Moving it from the mainframe to another server for post processing is not only time consuming but also expensive. More recently, a relatively new firm, Veristorm, Inc., of Santa Clara[46], has demonstrated mainframe capabilities for Big Data analysis most closely associated with *Hadoop Map Reduce* and huge arrays of x86 servers and RAID storage. This company is addressing the need for Big Data analysis of mainframe data by an offering exclusively for Linux on a System z called *VStorm Enterprise*. It includes a set of data collectors that facilitate the movement and transformation of z/OS data to *HDFS (Hadoop Distributed Files System)* files. It also includes *zDoop*, the first commercially-supported Hadoop distribution for System z.

Why do this? First, very large amounts of sensitive data do not have to be moved to distributed platforms, which is very costly and time consuming when done frequently. Second, in doing so, all the redundancies of servers, storage media, and data associated with HDFS systems are avoided. Third, it is very efficient. In recent tests at IBM's Poughkeepsie Laboratory, two billion records were collected and processed in two hours using the capacity of only two IFLs (roughly 2500 MIPS) and DS8800 storage for the transformed files. This suggests that new uses for the modern mainframe have yet to be limited.

Clearly, cloud thinking was driving much of the new innovation with the current mainframe offerings from IBM. If you truly want a hybrid cloud (say running Linux on mainframe IFLs, on Power blades and on System x blades, plus Windows on System x blades, with tightly-coupled analytics via IDAA and Veristorm, no one besides IBM is giving you that kind of open cloud diversity. Plus, if you have enterprise applications running on mainframes (say, on z/OS), especially with DB2 or IMS databases, you have many possibilities of tightly linking the open clouds with mainframe-resident data, without need to create and transfer one or more copies of that data to another location. This is enterprise cloud computing at the ultimate extreme.

(7) The Mainframe Focuses on Enterprise Clouds

Many of the recent improvements to the mainframe environment – some hardware and some software – have been focused on delivering enterprise-class cloud solutions. These include:

- Cross-platform coherence and management
- The Enterprise Linux Server Offering
- The Enterprise Cloud System Offering

Collectively, these make for well-rounded and well-managed mainframe-based cloud solutions. Each of these will be discussed.

Smarter Infrastructure Brings Coherence to Cross-Platform Management

Public and private enterprises can be said to share the common goal of more effective business results through (1) improving customer satisfaction through better reach and relationships; (2) becoming more agile to respond to rapid changes in their service environment: and (3) by infrastructure optimizing to reduce TCO, which translates into better value for IT investments. This needs to be accomplished by meeting the challenges of systems management processes that bridge the needs of both the system(s) of record (home to transaction, database, and command and control systems) and systems of engagement (where the interaction and collaboration with end users takes place).

Systems of record are commonly associated with System z mainframes. Systems of Engagement usually provide the on-ramp for diverse cloud workloads. Quite often, these activities have to be

[45] For more detail on the IDAA and analytics on zEnterprise, see **The Clipper Group Navigator** entitled *Addressing New Business Analytics Challenges - When the IBM zEnterprise Really Makes Sense*, dated December 21, 2012, and available at http://www.clipper.com/research/TCG2012030.pdf.

[46] See http://www.veristorm.com/.

accomplished with the highest qualities of service for security, reliability, and consistency – the qualities best demonstrated by System z's well-integrated hardware and software technology. The aforementioned bridging can be broken down into three related areas – *automation*, *performance management*, and *IT analytics*. Each of these areas is being addressed by a family of products from the IBM Tivoli portfolio. These include IBM's *Tivoli System Automation* family – *System Automation Application Manager*, *System Automation for z/OS*, and *System Automation for Multiplatforms*; *IBM Cloud Management Suite for System z*, the *IBM Tivoli OMEGAMON Performance Management Suite (OMEGAMON XE for z/OS, OMEGAMON XE for z/VM* and *OMEGAMON XE for Linux on System z*.[47]

Virtualization Management Enhanced by IBM Wave for z/VM

IBM's acquisition of *CSL International* in August 2013 brought the *CSL-Wave* product technology into the z/VM portfolio. Now called *IBM Wave for z/VM*, it offers very significant and innovative improvements to the management of larger and more complex deployments of z/VM virtual machines, especially Linux images. IBM Wave for z/VM organizes, unifies, and simplifies many operations that will result in significant productivity savings. Essentially, it is Linux-administrator friendly, in that it doesn't require traditional mainframe knowledge and experience and delivers a click, drag, and drop administrative console. It offers the following benefits.

- A single dashboard simplifies administrative and management of virtualized servers.
- Reduces time to perform complex management tasks.
- Extends the reach of existing (non-mainframe) skills.
- Provides accurate views of the system that result in improved quality and consistency of operations.
- Reduces risk of errors by delegating key tasks to appropriate authority.
- Accelerates steps required for virtual server cloning and provisioning.

There is no question that the combination of z/VM, IBM Wave, and Linux on System z in IBM's enhanced *ELS* solution offering (described in the next section) provides an ideal integrated virtualized infrastructure and virtualization management essential to a mainframe-based cloud solution.

The Enterprise Linux Server Reintroduced

The concept of an *Enterprise Linux Server (ELS)* on System z is not new. The first models of a System z that were offered with IFL-only configurations date back to January 2002, as discussed earlier.[48] Currently, IFLs account for one-third of System z MIPS shipped, and one-third of all System z clients have enabled Linux through this facility. The 2014 announcement of an entry level zBC12-based[49] ELS brings significantly greater scale and improved price-performance to this solution, compared to previous offerings. Up to 13 processor cores may be configured as IFLs supporting Linux for System z. Up to 40 virtual servers per processor can be deployed yielding up to 520 Linux server images in a single footprint[50]. The integrity of the Linux servers is protected by the exceptional reliability of System z processors and the secure isolation provided by z/VM's virtual server partitioning to provide superior levels of availability and security.

Red Hat and *SUSE Linux*, and *OpenStack Cloud* distributions are supported.[51] IBM is offering prepackaged Enterprise Linux Server solutions that includes all necessary processor hardware, hypervisor, and maintenance and comes with a total cost of ownership that promises to be very attractive, both operationally and with respect to ROI, when compared to x86-based distributed servers with a typical software stack.

In laboratory studies by IBM, a deployment of 3200 virtual machines on a zEC12 would result in cost savings as much 80% for energy, 90% for floor space, and 70% in labor costs over a three year period, all at 90% resource utilization rate. And, at the end of the period, the server assets can be upgraded at costs much less than total replacement of hardware and software typical for replacement of aging x86-based distributed servers.

---

[47] For more information on these software components, see the April 8, 2013 issue of **The Clipper Group Navigator** entitled *IBM zEnterprise is Enterprise Cloud Infrastructure*, and available at www.clipper.com/research/TCG2014008.pdf.

[48] The first IFL-only mainframe was the *z800 Model 0LF*.

[49] An ELS solution may include a either zEC12, zBC12, z196, or a z114.

[50] The actual number of images achievable is workload and configuration dependent.

[51] The Linux stack must be obtained from a qualified distributor, who in turn provides support for the Linux environment.

The IBM Enterprise Cloud System – Complete, Easy to Deploy, and Price Competitive

The *IBM Enterprise Cloud System* solution announced this week addresses the needs of enterprises and service providers that desire a low-cost, low-risk solution for entry into a scalable cloud deployment model. The ECS is a high-availability and secure solution that consists of a pre-configured, integrated and factory-tested system, including a processor (zBC12 or zEC12), networking, and storage (*DS8870* or *Storwize V7000*) hardware, with options for number of IFLs, memory, and I/O ports. The software ensemble includes z/VM 6.3 with several required features, IBM Wave, and the Cloud Management Suite for System z, which includes OMEGAMON XE for z/VM and Linux, *Tivoli Storage Manager*, and the *SmartCloud Orchestrator*. The IBM promise is for a simplified, bottom-line pricing model that will compete with an equivalent x86 configuration and can scale much more smoothly.

The Enterprise Cloud System is a broader offering than the Enterprise Linux Server and is more focused on providing a complete cloud solution. The ECS solution, though limited to the most current zEnterprise system processors, also includes a choice of storage and I/O configurations, and a more complete suite of software performance and management tools focused on the management of an enterprise cloud deployment. Of course, ELS configurations can, on an a-la-carte basis, include any software or hardware products or features that are included with the ECS.

Conclusion

We have travelled through fifty years of IBM's mainframe history with three intertwined threads.

- Historical highlights of the evolution of the IBM mainframe over its first 50 years.
- The concept of "cloud" which, while relatively modern terminology, has been part of the mainframe's persona for more than four decades.
- The focus of my career in mainframes with IBM and Clipper, which has paralleled these 50 years.

I hope that you enjoyed reading this as much as I enjoyed sharing my story and IBM's mainframe saga with you.

Here is the important take away. **IBM's zEnterprise offers an extraordinary set of enterprise cloud solutions.** Don't let popular thinking sway you from considering the most complete and contemporary enterprise cloud solution. Whatever the many challenges that enterprise datacenters will have in the next 50 years, I am convinced that IBM's mainframe technologies will be there to meet them. The first 50 years tell me so.

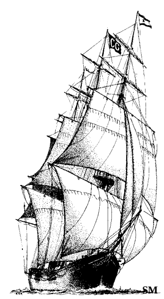

About The Clipper Group, Inc.

The Clipper Group, Inc., now in its twenty-second 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

***Stephen D. (Steve) Bartlett* is a Senior Contributing Analyst for The Clipper Group.** Mr. Bartlett's interests include enterprise solutions including servers, system software, middleware, their underlying technologies, and their optimal deployment for responsiveness to emerging business requirements. In 2010, he joined the Clipper team after over 42 years with the IBM Corporation as an account and program manager in large system sales, product development, strategy, marketing, market research, and finance. During that time, he received several awards for his contributions in innovative market research and contributions to the business. He has a B.S. from Rensselaer Polytechnic Institute, and an M.S. from Union College.

- ***Reach Steve Bartlett via e-mail at steve.bartlett@clipper.com or at 845-452-4111.***

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.