

The IBM Mainframe is Still the Real Deal

Keith Allingham

DataKinetics' CEO Allan Zander's article appeared on the [Planet Mainframe blog in 2021](#), and it made some waves in the distributed systems world—see [Hacker News discussion](#). The IT folks without much knowledge of IBM mainframe architecture refuse to believe the numbers in the article, while others lambaste the article for being old. Some of the comments are fair, but as AWS, Google, et al, provide more powerful cloud/distributed solutions, IBM does the same with its mainframe platform.

About 12 months ago, I had a conversation with an IBM tech in their computing costs group responsible for analysis and cost comparison between existing IBM Z installations against proposed cloud replacement proposals given to IBM customers by cloud business organizations (AWS, Google, Microsoft, etc.). He was interested in the Planet Mainframe article because it closely mirrored the results that he and his team were seeing in the 2020-2021 time frame, as they researched and compared an IBM Z installation vs an AWS proposal.

As you can see, this article cites data that is now more than 5 years old, but the comparison is still generally accepted as valid. Until a new article with new data is available, here is the Planet Mainframe article that started all the fuss:

The IBM Mainframe: The most powerful and cost-effective computing platform for business

Allan Zander, CEO, DataKinetics

Many of we mainframe pundits have written about the robustness, power, perseverance, capacity and more importantly, the cost-effectiveness of the mainframe ([Allingham](#), [Sun](#), [Peleg](#)), including [myself](#). But what about showing the superiority of the mainframe using real numbers, comparing it to other platforms? That requires a lot more work. [Schroder](#) and [Olders](#) show us some real-world numbers, but how about showing the ugly details? That's even more work, and fortunately, a couple of people have done that as well.

Michael Benson's Enterprise Executive article in 2016 did that—since then, distributed servers have come a long way (AWS, Google and a host of other cloud service providers), but so has the mainframe. In 2015, the top-of-the-line mainframe was the z13, an outstanding business machine; since then, successive machines, z14 through z16 (and counting) outperform it considerably on many levels –speed, transaction throughput, security, flexibility, and more. A main argument then, as now, is cost; and that's a losing argument right from the get-go.

Comparing Platform Costs

“Other platforms are cheaper...” This is the basic claim for most people interested in dumping mainframe systems in favor of commodity servers. The argument is simple: “Google, Amazon and Microsoft don't use mainframe systems at their back end, so why should anyone?” Fair question, but let's look at the premise first—are server farms less costly than the mainframe? Recently, Michael Benson did a study for Enterprise Executive magazine in an article called CIOs: Are You Really Paying Less by Using x86 Platforms? In it, he configured two similar performing platforms—one mainframe-based, using an IBM z13 mainframe system, and the other, a bank of HP servers. Table 1 shows the system specifications.

He explains that running Linux on the mainframe is no different than running it on x86 servers. The only real difference is the cost, and the belief is that x86 platforms do it for less. But do they? The hardware costs for these configurations run in at \$2,299,451.00 for the server farm solution, and \$2,793,371.00 for the mainframe solution. However, due to licensing costs, the software cost for the server farms comes in at \$1,807,406.00, with the mainframe running at only \$416,883.00.

So yes, the hardware is cheaper, but there is not quite as much difference as you might expect. The real surprise is the difference in software cost. When you also consider maintenance costs, the pattern continues. Maintenance

ATTRIBUTE	HP PROLIANT BL460 GEN9	IBM Z13 2964 N30
Total Servers	12	1
Processors	24	30
Cores/processor	12	1
Cores/server	24	30
Total cores	288	30
# VMs	1000	1000
Memory	2 TB	2 TB
Hypervisor	VMware vSphere 4	IBM z/VM
Cloud Mgmt	VMware vRealize	IBM Wave
OS	Red Hat Enterprise Linux	Red Hat Enterprise Linux
Web server	Apache HTTP	Apache HTTP
Application server	IBM Websphere	IBM Websphere
Messaging	MQ	MQ
Database	Oracle EE	Oracle EE

ANNUAL LABOR	HP PROLIANT BL460 GEN9 (QUANTITY=12)	IBM Z13 2964 N30 (QUANTITY=1)
Server admin	\$580,160	\$430,000
Net admin	\$384,000	\$76,800
Total	\$964,160	\$506,800

costs for the server farm come in at \$390,327.00, with the mainframe at \$269,767.00. Labor costs are also part of the picture.

At the end of the day, what really matters is the ongoing operational costs of the two platform solutions. Table 3 shows a considerable gap in favor of mainframe computing.

Over a five year period, operating costs compound, and the picture looks much worse for the server farm, \$9,052,749.00 vs \$6,979,693.00 in for the mainframe setup. The shocking conclusion therefore, is that it is cheaper to run the mainframe system than it is to run the server farm.

OPEX	HP PROLIANT BL460 GEN9	IBM Z13 2964 N30
Hardware mtce	\$9,544	n/a
Software mtce	\$390,327	\$269,767
Admin	\$964,160	\$505,800
Other (power, etc.)	\$31,505	\$68,355
Total	\$1,395,536	\$844,922

When doing cost comparisons, it is good practice to look at all contributing costs, and to look at long term cost of ownership. This comparison would have looked a lot different if we stuck to just the hardware acquisition cost, or even if we hid the personnel costs in a general employee pool rather than in the TCO calculations.

Technology Economics

Cost is one thing—often a very misunderstood thing, as Michael Benson pointed out. But acquisition and ongoing cost represent only one dimension in a complicated cost-comparison between platforms. What about environments that run a mix of mainframe and distributed systems? And what about comparing not just cost between platforms, but real costs in specific industries? Well, that’s where Dr. Howard A. Rubin of *Rubin Worldwide*, a technology economics research firm, comes in.

In his paper, *The Surprising Technology Economics of Mainframe vs. Distributed Servers*, Dr. Rubin explains that understanding computing platforms and their economic relevance in the context of their contributions to business performance is critical. This context provides a transparency that goes far beyond the basic economics of the costs of hardware and software acquisition or a TCO calculation. This is especially important when we consider that technology costs are a rising part of ongoing business operations expense.

It is good practice to look at all contributing costs, and to look at long term cost of ownership.

IT costs vs business revenue and cost

Technology costs relative to business revenue and operating costs vary considerably from one industry vertical to another. For example, in banking and finance, IT expense represents about 6% of revenue and just over 7% of business operating expense; compared to the retail sector, where IT expense represents just under 1.5% of revenue and just over 1.5% of business operating expense.

Cost of platform choice

Businesses have choices on how they will handle their processing needs and this typically comes down to the mainframe and server farms. The cloud is part of the latter solution. The reality is that any business that runs mainframe systems also runs server farms, so it is fair to characterize them as running “mainframe-heavy” datacenters, while those without mainframe run “server-heavy” datacenters. It is also useful to consider new metrics for these datacenters—MIPS per \$1M of revenue and physical servers per \$1M of revenue. These aren’t equivalent in any way, but they serve to represent the economics of their computing choices in measurable economic terms.

When comparing businesses within the same industry vertical, the “heaviness” of their IT deployment strategies result in a significant differences. For example, for financial services businesses:

Mainframe-heavy shops consume:

- 3.1 MIPS per \$1M of revenue
- 0.22 servers per \$1M of revenue

While the server-heavy shops consume:

- 1.75 MIPS per \$1M of revenue
- 1.2 servers per \$1M of revenue

BUSINESS	AVERAGE COST	MAINFRAME-HEAVY COST	SERVER-HEAVY COST
Distribution	\$4,255,273	\$3,936,728	\$6,809,818
Communications	\$4,979,371	\$4,306,000	\$8,295,000
General business	\$4,832,000	\$4,414,000	\$7,846,000
Computer Services	\$6,093,958	\$5,644,350	\$7,619,000
Industrial	\$9,270,513	\$9,082,000	\$11,181,000
Financial Services	\$12,627,002	\$12,742,000	\$16,445,000
Government	\$15,161,129	\$14,148,000	\$15,981,703
Average	\$8,174,178	\$7,753,297	\$10,596,789

When these figures are mapped to the total cost of mainframe and server farm costs within various industry verticals, the economic differences that can be attributed to their deployment strategies become apparent—(Table 4). The inescapable conclusion is that mainframe-heavy computational costs to support a \$1B organization on average may be 30% lower than a server-heavy deployment.

Cost of Goods

While the cost of technology yields interesting conclusions, the actual costs of platform choice are also surprising, and support the former. The next step is to link the technology costs to business costs.

A good way to do that is to use a cost-of-goods metric. Ask the question: “What is the IT cost contribution to the business cost of goods?” And follow that up with: “How does technology deployment affect the measure of impact on the business?” Table 5 itemizes the cost of goods for five business types—finance, industrial, communications, general business and insurance.

This data implies that where appropriate, effective use of mainframe resources results in a 29% cost advantage over distributed server-heavy deployments.

Looking closely at the insurance data, we see that the average IT cost of processing an insurance claim in a mainframe-heavy environment is approximately \$56, which is \$36 less than the processing cost in a server-heavy environment. What does that mean to an insurance business? For an insurer that processes 100,000 claims per year, the savings could be \$3.6 million per year by leveraging mainframe technology.

Similarly, a bank with 4500 ATMs would be paying over \$1000 per ATM using a server-heavy datacenter, as compared to less than \$600 using the mainframe-heavy scenario. Such a bank could save more than \$2 million per year by leveraging mainframe technology. To be fair of course, ATM costs are only one small part of a bank's IT cost concerns.

Competitive advantage

Any large company interested in maximizing computing power AND controlling costs will clearly enjoy a competitive advantage over a similar company that just seeks to avoid mainframe technology in favor of server farms. This advantage translates directly to the bottom line, shareholders and investors. And for a company considering a mainframe migration project as a means for cutting costs, this information could be seen as "found money."

PROCESSING COST PER:	AVERAGE COST	MAINFRAME-HEAVY COST	SERVER-HEAVY COST	RATIO, MF VS SERVER
ATM	\$928.00	\$572.00	\$1,021.00	56%
SKU	\$227.27	\$184.09	\$252.27	73%
Mobile subscriber	\$23.26	\$18.26	\$26.12	70%
Patent	\$390.83	\$372.00	\$401.00	93%
Claim	\$78.00	\$56.00	\$92.00	61%
			Average	71%

Conclusions

The facts support the notion that the mainframe is the most powerful and cost-effective computing platform for large businesses with a need for high-intensity transaction processing. Claims to the contrary are typically either as a result of simple lack of knowledge on the subject, or a biased unwillingness to look objectively at the facts.

But if the mainframe is so great, then why is it not being used by the newest and latest concerns (Amazon, eBay, etc.)? The reason is bias. Whether intentional or through ignorance, there is a great deal of bias against the mainframe. We hear it all the time – and saw it in the comments to the original publication of the article. People say "It's too expensive!" (It clearly is not.) "It's old and dusty!" (Obviously not.) "It's hopelessly outdated!" (Not even close.) "I don't know very much about it!" (Ahhhh . Now we're getting somewhere.)

The last part is the key to the puzzle of why the mainframe generally has a difficult time displacing server farms in environments where it could make a positive impact. The truth is, organizations that could benefit from the mainframe, but don't, are leaving money on the table.

So, What's the Strategy?

We've told you what IBM already knows, what many IBM customers already know, what some "mainframe replacement" vendors secretly know, and even what today's big cloud vendors know. So, what's the strategy moving forward?

What about mainframe shops having trouble keeping up with growing workloads on their "most powerful and cost-effective" mainframes? Should they be upgrading? Shifting workloads off-platform? As you might guess, there are options. There are a couple of organizations that are helping mainframe shops to optimize what they have now—to increase workload throughput of the systems they're currently running. No upgrade needed; no changes to application logic, no changes to the Db/2 (or whatever) database being used. This is possible using high-performance in-memory technology.

And both IBM and DataKinetics are offering these solutions right now. Talk to people who actually know something about the platforms under evaluation.



[Click Here to Watch the Video](#)