Technology Economics:

Economics of Computing -The Internal Combustion Mainframe [Expanded Version]

Dr. Howard Rubin
CEO and Founder Rubin Worldwide
howardarubin@gmail.com
www.rubinworldwide.com
+1 914 420 8568

For a few moments, don't even think about information technology (IT)…. think about the automobile. There are roughly 750 million to 900 million of them on our planet. And even though today we have hybrids, synergy drive, electric cars, and diesel, the internal combustion engine powers 99% of them and is still predicted to be the main engine of transportation with 50% penetration through at least 2025.

The internal combustion engine has been with us more than 100 years. And yes, there are major issues with its use of oil and the impact of its emissions on client change, but consider how it has evolved and been tuned and refined based on a century of experience.

In fact this became more than obvious to me on May 27th as I attended the U.S. introduction of the Ferrari 458 in Greenwich CT. There, under glass, as an integral part of this automotive masterpiece of art and design ,was a 275 cubic inch internal combustion engine producing 580 horsepower or about 2.11 horsepower per cubic inch. It was not quite my "father's Oldsmobile" (he actually had a 1970 Buick and yes, Oldsmobile is gone) – which had 350 cubic inches and produced 125 horsepower or .24 horsepower per cubic inch at likely less miles per gallon than the Ferrari. Perhaps Moore's Law works for engines too. In 40 years we have seen almost a 10x improvement in efficiency. And that improvement is just really hard to discount or discredit and toss all away in the name of replacing all computational engines with the "flavor of the day" – servers, blades, perhaps the "cloud".

By the way, please don't assume at this point that this article is about mainframe superiority. Rather it is about the economics of computational platforms and the criticality of transparency into this dimension of technology economics in maximizing the value of IT to your business.

Most businesses today rely on an un-engineered mix of computational platforms – mainframe computers; UNIX, Wintel, and Linux Servers; Midrange devices (e.g. AS400). The choice of how much to put on the mainframe or servers just sort of happens, it is not by design. At the same time, many organizations have adopted standards and architectures that move them away from mainframe computing in the near and far future under the banner of modernization. However, while mainframe computing (or server-based computing) may not be right for all forms of computation – notice I haven't even gotten to supercomputing and megaflop environments – it is essential for organizations to simultaneously consider both the functional characteristics of their computing needs and economic considerations.

An analysis of data from 21 sectors (inclusive of governments) and 133 companies across those sectors reveals that the average company has about .37 MIPS per $1M of revenue and about .17 servers.

In a more business-like context, this would mean that the average $10B company would have a core platform of 3,700 MIPS and 1,700 physical servers. (Obviously as virtualization takes hold this will change – sort of like with the introduction of fuel-injection and multi-valve and variable cam technology with electronic controls modules on the internal combustion engine.) However, all sectors are not created equal in terms of computational needs. Banking and finance might require 10,700 MIPS and 4,600 servers to support $10B in revenue on average. While at the lower computation load end of the spectrum a professional services company would require 1,400 MIPS and 800 servers.

Such platform demographics can be translated into "compute cost per $1M revenue" by using Gartner IT Key Metrics Data (2010) for the average MIPS and physical server operation cost per year -- $4,445 per MIPS and $10,473 per server respectively. The result of such an analysis indicates that a bank or financial serves company has a compute cost of about $9,574 per $1M revenue while the professional services company is at $1,460 [see Figure 1].

Figure 1

	Sector Averages n=133 companies				
	MIPS per $1M Revenue	**Servers per $1M Revenue**	**Cost per MIPS (Source: Gartner 2010 Key Metrics)**	**Cost per Server (Source: 2010 Gartner Key Metrics)**	**Compute Cost (MIPS + Server Cost) per $1M Revenue**
Banking	0.98	0.39	$ 4,445	$ 10,473	$ 8,441
Consumer Products	0.19	0.16	$ 4,445	$ 10,473	$ 2,520
Education	0.13	0.05	$ 4,445	$ 10,473	$ 1,102
Electronics	0.25	0.11	$ 4,445	$ 10,473	$ 2,263
Financial Services	1.07	0.46	$ 4,445	$ 10,473	$ 9,574
Food & Beverage Processing	0.18	0.12	$ 4,445	$ 10,473	$ 2,057
Government - Federal	0.49	0.12	$ 4,445	$ 10,473	$ 3,382
Government - State & Local	0.38	0.09	$ 4,445	$ 10,473	$ 2,632
Health Care	0.19	0.13	$ 4,445	$ 10,473	$ 2,206
Insurance	0.33	0.16	$ 4,445	$ 10,473	$ 3,143
Manufacturing	0.21	0.12	$ 4,445	$ 10,473	$ 2,190
Metals & Natural Resources	0.16	0.12	$ 4,445	$ 10,473	$ 1,968
Professional Services	0.14	0.08	$ 4,445	$ 10,473	$ 1,460
Telecommunications	0.85	0.25	$ 4,445	$ 10,473	$ 6,397
Transportation	0.23	0.21	$ 4,445	$ 10,473	$ 3,222
Utilities	0.16	0.08	$ 4,445	$ 10,473	$ 1,549
Cross Industry Average	0.37	0.17	$ 4,445	$ 10,473	$ 3,382

Other interesting data points for computer cost per $1M revenue by sector are: Insurance $3,143, Telecommunications $6,397, Healthcare $2,206, and Federal Government $3,382.

While this information is somewhat interesting, it does not get at the central issue regarding the economic tradeoffs between mainframe and server computing choices (of course with the underlying assumption of workload appropriateness).

On a sector by sector basis, an analysis of companies of similar size and business mix that exhibit polar opposites in computing platform choices – mainframe centric or server centric – provides insights into developing an understanding of the essential economics.

For example, in the study database there are numerous ~$10B banks. The platform demographics indicate at one end of the spectrum there is one operating with 10,700 mainframe MIPS and 4,600 servers and at the other end one is operating at 7,643 MIPS and 8,846 servers.

Using the aforementioned cost per MIPS and cost per server data indicates that the mainframe centric organization has an expense of $95.7M while the server centric organization in the same business has an expense of $126.6M in supporting the same business. Of course, there are numerous other considerations than such raw economics in making platform and architectural choices, but few organizations have taken the time examine the indicated tradeoff. And if you scale this up to a $100B enterprise, the spread is $300M – that is real money!

Extending this analysis across all 21 study sectors indicates that our "average $10B company" would use 3,700 MIPS and 1,700 servers as mentioned earlier at a cost of $33.8M [see Figure 2]. Its mainframe-heavy clone would use 4,500 MIPS and 850 servers at a cost of $29.1M while its server-biased comparator would use 2,600 MIPS and 3,180 servers at a cost of $45.1M (see www.rubinworldwide .com for a table-of the-elements with all this data). Our hypothetical company at the average is $4.7M higher in expense that the mainframe biased model. The server biased company is $15M or 33% higher than the mainframe pole. Likely, if this was transparent to business leadership the question would be asked – "What is the business value of this extra expense?" [Seems like sort of the issues we hear about today with power choices for automobiles. And as is with the automotive question, the answer lies in the context and not just with the numbers.]

Figure 2

				Mainframe Centric Model				Server Centric Model			
	Platform Size and Costs to Support $10B			Platform Size and Costs to Support $10B Revenue				Platform Size and Costs to Support $10B			
	MIPS	Servers	Platform Cost at Gartner 2010 Benchmark Unit Costs	MIPS	Servers	Platform Cost at Gartner 2010 Benchmark Unit Costs	Gap to "Average Model"	MIPS	Servers	Platform Cost at Gartner 2010 Benchmark Unit Costs	Gap to Average Model
Banking	9,800	3,900	84,405,700	12,000	2,000	$ 74,286,000	$(10,119,700)	7,000	7,500	$ 109,662,500	$ 25,256,800
Consumer Products	1,900	1,600	25,202,300	2,327	821	$ 18,934,659	$ (6,267,641)	1,357	3,077	$ 38,257,115	$ 13,054,815
Education	1,300	500	11,015,000	1,592	256	$ 9,761,099	$ (1,253,901)	929	962	$ 14,197,692	$ 3,182,692
Electronics	2,500	1,100	22,632,800	3,061	564	$ 19,514,989	$ (3,117,811)	1,786	2,115	$ 30,091,923	$ 7,459,123
Financial Services	10,700	4,600	95,737,300	13,102	2,359	$ 82,944,110	$(12,793,190)	7,643	8,846	$ 126,618,269	$ 30,880,969
Food & Beverage Processing	1,800	1,200	20,568,600	2,204	615	$ 16,242,066	$ (4,326,534)	1,286	2,308	$ 29,883,462	$ 9,314,862
Government - Federal	4,900	1,150	33,824,450	6,000	590	$ 32,846,385	$ (978,065)	3,500	2,212	$ 38,718,942	$ 4,894,492
Government - State & Local	3,800	900	26,316,700	4,653	462	$ 25,516,549	$ (800,151)	2,714	1,731	$ 30,191,346	$ 3,874,646
Health Care	1,900	1,300	22,060,400	2,327	667	$ 17,323,429	$ (4,736,971)	1,357	2,500	$ 32,215,000	$ 10,154,600
Insurance	3,300	1,600	31,425,300	4,041	821	$ 26,554,659	$ (4,870,641)	2,357	3,077	$ 42,702,115	$ 11,276,815
Manufacturing	2,100	1,200	21,902,100	2,571	615	$ 17,874,923	$ (4,027,177)	1,500	2,308	$ 30,835,962	$ 8,933,862
Metals & Natural Resources	1,600	1,200	19,679,600	1,959	615	$ 15,153,495	$ (4,526,105)	1,143	2,308	$ 29,248,462	$ 9,568,862
Professional Services	1,400	800	14,601,400	1,714	410	$ 11,916,615	$ (2,684,785)	1,000	1,538	$ 20,557,308	$ 5,955,908
Telecommunications	8,500	2,500	63,965,000	10,408	1,282	$ 59,691,209	$ (4,273,791)	6,071	4,808	$ 77,338,462	$ 13,373,462
Transportation	2,300	2,100	32,216,800	2,816	1,077	$ 23,797,187	$ (8,419,613)	1,643	4,038	$ 49,597,308	$ 17,380,508
Utilities	1,600	800	15,490,400	1,959	410	$ 13,005,187	$ (2,485,213)	1,143	1,538	$ 21,192,308	$ 5,701,908
Cross Industry Average	3,713	1,653	33,815,241	4,546	848	$ 29,085,160	$ (4,730,081)	2,652	3,179	$ 45,081,761	$ 11,266,520
							-14.0%				33.3%

Bringing this analysis down one level to IT cost of goods [see Figure 3], our cost per trade would move down $.013 with an optimized computational platform of mainframe server mix – multiply that by a few billion and it gets interesting. Healthcare costs are a major issue today. Before the injection of $20B into healthcare IT by the US administration, the IT cost per hospital bed (947,000 of them in the US; whether occupied or unoccupied) is $65 per day. The optimum computational mix would drop that by $4-$6.

Figure 3

Industry	Measure	Average IT Cost of Goods	Mainframe Biased	Server Biased
Airlines	Per Passenger Mile	$ 0.007	$ 0.0061	$ 0.0076
Automotive	Per Vehicle	$ 333	$ 275	$ 370
Chemicals	Per Patent	$ 57,717	$ 55,800	$ 59,552
Consulting	Per Consultant	$ 53,060	$ 48,900	$ 62,344
Hospitals	Per Bed per Day	$ 64.30	$ 54.4000	$ 71.7000
Railroads	Per Ton Mile	$ 0.0014	$ 0.0012	$ 0.0018
Retail	Per Store (Door)	$ 494,818	$ 421,346	$ 560,300
Web Sites	Per Search	$ 0.042	$ 0.046	$ 0.041
Trucking	Per Road Mile	$ 0.177	$ 0.1550	$ 0.1940
Armed Service	Per Person	$ 8,036.00	$ 6,871.00	$ 9,839
Utilities	Per MegaWatt Hour	$ 2.63	$ 2.21	$ 2.94
Oil & Gas	Per Barrel of Oil	$ 2.10	$ 1.78	$ 2.32

Technology economics is clearly still in its infancy. However, with 50 years of history or more, the mainframe itself is mature and reliable as the internal combustion engine and still evolving. Today's newer platforms – servers, blades, cloud, etc. – will evolve themselves and clearly be under pressure as technologies that are barely envisioned today enter the marketplace. Decisions about computing need to be made in the context of relevance. Relevance implies applicability to the business application on hand and a true understanding of economics in terms of expense and real value.

Interestingly enough, the simplistic analysis described herein reveals more of what we don't know than what we do know. It is clearly time to explicitly engineer the economics of core computing platforms with the little that we do know…and the little that we do know now is that the marketplace doesn't reward firms that are using only the latest technology at any expense, in the long run it rewards those that make the optimum use of the right computing resources in the right way as evidenced by business performance.