

Strategic Advisory

Refresh considerations for data center architecture (Intel/AMD-based racks and blades)

AUGUST 2011

Need

Large enterprises need data center architectures which are stable, efficient, cost effective, and matched to the demands of their internal customers. Providers of cloud services have the same requirements, but have the additional requirement of being competitive with other providers. This competition can be based on cost, quality of service, specialization, or a mix of these. For both large enterprises and cloud providers, decisions about refresh cycles and upgrades must factor in costs and performance demands.

Goal

This discussion attempts to identify several major quantitative and/or financial factors relevant to making

refresh/lifecycle and financing decisions about Intel-based data center systems. The goal is to set forth some base level evaluation models for use in calculating the effects of technology changes on financial and system performance metrics.

Assumptions

This discussion assumes that the data center consists of high performance rack servers or blade servers, populated with systems composed of Intel Xeon (and AMD equivalents) processors and (optionally) on-board hard disk drives. It further assumes that most units are replaceable in-line ('hot swappable'), and that system management and/or provisioning software can redistribute workloads within select pools of servers.

Approach

There are two steps in this approach:

- estimate the impact of shorter refresh cycles on real costs
- estimate the impact of leasing on the investments required to move to shorter cycles.

Step 1 – Estimation of the financial impact of shorter equipment refresh cycles

Energy costs

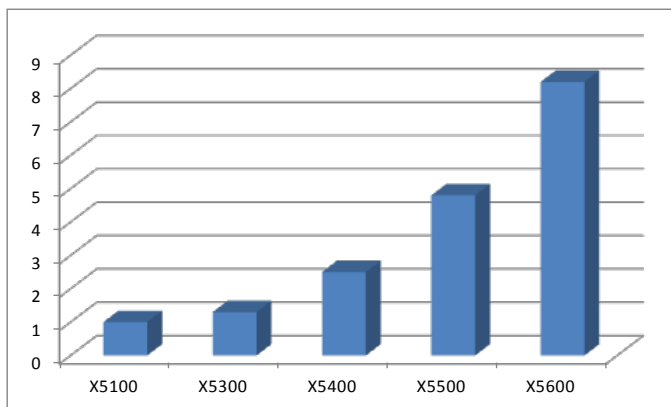
A growing reason to move to shorter refresh cycles today (especially for volume servers) is the steady advance of energy-efficiency technology. The servers available today use only a fraction of the energy used by prior generations of servers, and this can amount to significant savings. Although we have only recently begun to measure performance per watt, all the available data indicated that each successive generation of Intel/AMD servers produces the same computational output for only 60 percent (and sometimes only 25 percent) of energy expense.

For example, prior to the creation of the SPECpower_ssj2008 benchmark, one test showed that the Intel Xeon 5160 performed 1.85x the workload per watt as did the prior-generation Xeon 5070. The SPECpower_ssj2008 benchmarks have shown similar results for advances within the Xeon family.

To help size this for financial assessment, if we use the Xeon 5160 (2006 technology) as a base of 1x, then the generation of Xeons which large enterprises installed in 2008 would do 2.45x the work per watt, the newer/2009 Nehalem Xeon 5500s would do 4.8x, and the latest 5600s would do 8.2x. By estimating replacement footprints, we could calculate energy savings.

Server performance per watt (relative to x5100)

*based on SPECpower measurements @ spec.org



The estimation process would be fairly simple: build a spreadsheet with approximate numbers of each processor 'layer' (eg Xeon 51xx, 53xx, 54xx) and apply the above ratios to calculate a relative energy savings. The watts per system could be taken from vendor spec pages, or more accurately, from representative systems on the SPEC-power website. [Note that this is both relative and 'rough', in that it does not isolate HDD energy, power supply efficiency, etc.]

We should also note that the vendors improve upon efficiency each year (with greater reductions in some years than others) with the implication that energy savings could increase every year, if the servers are consistently being refreshed.

Here's the model run for 1,000 units of four generations of Xeon chips, with annual kW dollar savings (at 100 percent utilization):

	Units	Watts	Extended kW	5600 equiv.	#5600s needed	5600 kW	Freed up watts	kW costs	Saved Ext kW \$
x5100	1,000	258	258	0.12	123	29	229	\$0.08	\$160,169
x5300	1,000	334	334	0.16	158	38	296	\$0.08	\$207,445
x5400	1,000	276	276	0.30	301	72	204	\$0.08	\$142,860
x5500	1,000	237	237	0.59	590	142	95	\$0.08	\$66,825
Total									\$577,300

Since each watt of power used for a server also requires one to two watts of power to cool the server, the estimated energy savings should be at least doubled (for the HVAC operational cost, excluding the equipment cost of the HVAC units).

Harvesting these savings, however, will typically not be under the control of IT, since utility bills are more frequently placed in the facilities and/or physical plant budgets and not in IT.

This means that IT would see the costs for the new equipment, but not see any reduction in its own budget (since utilities are not often a line item in the IT budget). The CFO therefore must adjust the budgets of facilities/physical plant to harvest these cash reductions in utility bills. In the cases where IT does have a budget line item for utilities (eg in dedicated, stand alone data centers), the savings can be harvested within the IT budget.


Cost avoidance/consolidation savings

We have already seen how performance per watt can reduce energy bills, but performance per dollar can also be used to reduce equipment expenses while keeping computing power at the same level. This is a simple reality of incessant technology improvements by vendors.

There are many ways to measure this, but just to use a common benchmark to size this (TPC-C), here's a chart of the computing performance/price for Intel servers over the past few years:

How much x86 computing power does a dollar buy? (Relative to 2001)

(based on TPC-C benchmarks @ tpc.org)



What this basically means is that a server today (from the 2010 column) can do the same amount of work per dollar as ten servers from 2001, or two servers from 2008. In other words (all other things being equal) a manager could recover half the space, half the support labor, and at least half the risk of component failure, by replacing servers every 18-24 months or so.

This growth in 'performance per dollar' allows us to either buy more computing power for the same budget, or buy the same computing power for less budget, or a little of both.

To estimate the savings from displacing these servers, one would need to estimate the displaceable 'carrying costs' of these units (eg cost of real estate, per unit labor/support costs, cabling costs, and other related asset costs such as insurance, security), and apply these to estimates of the installed base (again layered by time). The cost calculation needs to be a 'net reduction' calculation, accounting for the new 'carrying costs' associated with the new units.

Here's the model run for 1,000 units of various ages/years of life:

	Units	Carry Cost	Extended carry \$	5600 equiv.	#5600s needed	Net unit reduc.	Extended net reduc.
In year 5	1,000	\$250	\$250,000	0.12	123	877	\$219,325
In year 4	1,000	\$250	\$250,000	0.16	158	842	\$210,429
In year 3	1,000	\$250	\$250,000	0.30	301	699	\$174,847
In year 2	1,000	\$250	\$250,000	0.59	590	410	\$102,454
						Total	\$707,055

Drive failures



By far, the biggest operational issue for this type of data center is the pattern of on-board hard drive failures. Numerous studies in the 2007-2008 timeframe (eg by Carnegie Mellon University, Google, University of Illinois) documented that HDDs have both a high 'infant mortality' (failure rate in the first three months for high usage systems) and a high failure rate in years three to five.

Here are the charts from the Google study:

Annualized failure rates by year and usage

('Failure trends in a large disk drive population')

Google study, p. 4-5.



If 6-8 percent of all HDDs are going to fail in the third year, and another 4-6 percent in the fourth year, then it is a simple matter to calculate how many replacements should be expected, based on the age of the equipment.

Calculating the cost per replacement has two variants: a failure within the warranty period and a failure after warranty expiration. Depending on how service is performed on the units (eg self-maintenance, substitute and ship), the cost will always have a labor cost, but the replacement equipment cost can be zero (under warranty) or non-zero (for out of warranty).

An argument can be made (and the impact calculated somewhat) that it is more cost-effective to refresh the systems (motherboard and storage together) some time during year three, before the AFR increases dramatically. If such a refresh program is implemented, then the unit is only touched once at decommissioning, prior to the HDD failure.

If, on the other hand, the HDD is replaced upon failure in years three or four (one touch), the unit will still have to be touched a second time when some other component eventually fails (eg ports, motherboard, memory, power supply) later. Since at least one of these other components will fail by the end of year five, a fix-in-place strategy will double the labor costs involved in maintaining the units (two touches versus one touch).

A rough way to estimate the cost range of this is to build a spreadsheet of the 'layers' of technology in the data center based on age. In the absence of hard data, management could simply start with known/approximate unit counts from a recent year and extrapolate backwards. Then, using the failure rates from the above charts, the predicted failures could be timelined and labor costs estimated.

For the units expected to fail within the five year planning horizon, labor costs (and equipment costs, if the

failure is expected after the warranty period) should be calculated and this number would be the labor savings from moving to a 'within three years refresh'. [In older gear, replacement HDDs may not actually be available, or price feasible.]

Then, the replacement equipment costs (for those units outside of warranty) would be added to the labor figure, to get a ballpark idea of the additional costs incurred by not moving to a three year refresh cycle.

See the chart on the next page which depicts the model run for 1,000 units in each year. Note, since the installed base likely contains units from each year/row in this table, the annual figure would be the total of the five rows.

Notice that these costs are bare bones for drives only (we logically would need to do one version of this table for all types of failures in rack servers: power supply, ports, memory, NIC), and also do not represent the significantly greater cost of business disruption.

	Units	Failure%	Failures	Labor cost	HDD\$	Equip\$	Extended
In year 5	1,000	7.00%	70	\$250	\$250	\$17,500	\$35,000
In year 4	1,000	5.90%	59	\$250	\$250	\$14,750	\$29,500
In year 3	1,000	8.20%	82	\$250			\$20,500
In year 2	1,000	7.80%	78	\$250			\$19,500
In year 2	1,000	3.00%	30	\$250			\$ 7,500
Total							\$112,000

Software licenses

There are significant savings possible (through basic consolidation) from the reduction of software and software support licenses. Although not strictly a linear function, consolidating ten servers into one also generally reduces the number of software/support licenses from ten to one. For ten servers, the server OS and utility software can easily represent \$5,000 to \$10,000 in annual contracts, and a reduction to one-tenth of this would net savings around \$4,500 to 9,000 for every ten servers consolidated.

In our hypothetical case, with a net reduction of over 2,500 servers (above under consolidation savings), the savings could be considerable: a net of 2.5K servers displaced, at \$500 annual software maintenance fees, could yield at least \$1.25m savings per year.

This reduction in software license and support fees is not present in situations where the consolidation is done by means of typical virtualization technologies. When we convert ten physical machines (PMs) into ten virtual machines (VMs) and run them on one larger physical server, we still have to pay for the software licenses inside the individual VMs plus the software for the larger host server (OS, utilities, management software, etc). Often, software vendor pricing can be reduced through negotiations and/or ‘virtual use’ pricing tiers, but these reductions are much less than those under traditional server consolidation. [The higher costs for virtualization are due to its higher value: we often use virtualization as

a tool to enable consolidation, to increase server utilization, to increase manageability and reliability, and other important tasks.]

Often a data center will ‘mix and match’ consolidation approaches—some virtualized, some not. So, we might very conservatively estimate our savings in this category at slightly more than a third of what we could achieve without virtualization: \$450,000.

Other issues

There are many other cost factors which should be included, mostly from the TCO studies (eg cost of maintaining multiple BIOS levels, multiple patch levels, wasting network ports by older servers, the move to SSDs) and there are many architectural issues which should be considered (eg ratio of virtual machines to physical machines, use of auto-provisioning and load-balancing systems to reduce refresh labor costs, higher density to reduce back-end decommissioning costs). The main issues discussed above would provide the baseline metrics for analyzing the financial aspects of (and needs for) a regular refresh of an Intel-based data center server infrastructure.

Assessments

Apart from other benefits, the estimates above make a strong case for a shorter refresh cycle for data center Intel/AMD servers. A quick table of the possible savings shows this clearly:

Savings category	Savings amount
Energy/utilities	\$577,300
Consolidation	\$707,055
HDD failures	\$112,000
Software licenses	\$450,000
TOTAL	\$1,846,355

Step 2—Estimation of the financial impact of leasing

Leasing and shorter refresh cycles: 'cash out the door' savings

All of the above savings argue for (1) faster refresh and/or (2) server replacement/consolidation, but all of these require either replacing existing equipment or implementing a faster refresh on all future server acquisitions.

In either case, leasing will prove to be less cash-expensive than purchase and this cost saving is incremental to the above savings. There are other operational and financial advantages of leasing, of course, but even a simple buy versus lease payment comparison illustrates the cost advantage clearly. [Note that there are ways of flexing these numbers a couple of points in either direction.] Using \$1,000 for a base comparison amount, here are typical cash out the door savings for the various refresh terms:

Months	Lease Pymts	Savings	Savings %
12	\$721.80	\$278.20	28%
18	\$808.11	\$191.89	19%
24	\$862.80	\$137.20	14%
36	\$966.60	\$ 33.40	3%
42	\$1,050.42	\$ (50.42)	-5%

of distributed servers will provide the opportunity of more savings in energy, licenses, repair, and footprint. Large enterprises and cloud service providers who run their businesses upon these architectures can exploit these savings for value for their stakeholders (customers, owners, employees).

When combined with various types of leasing structures from Macquarie Equipment Finance, even more savings can be achieved, while implementing improved asset management practices and insuring that future cost savings are more easily harvested.

This represents savings on equipment invoices, over and above the savings from refresh/consolidation. As a means to actually lower the equipment cost (the 'I' in ROI), this will accordingly raise the return (the 'R' of ROI) as well.

This element is easy to estimate: simply apply the percentage savings above to the purchase price of new servers for the appropriate term.

This means that to the extent the data center implements a shorter refresh cycle (to reduce costs), to that same extent, leasing will prove to be the fiscally preferred financing vehicle (to reduce further costs).

Leasing and other considerations

The leasing payment model above is a simple, easy to understand model, but there may be other operational and financial reasons to lease, apart from equipment or technology-based reasons. These other factors might involve strategic use of capital, improved asset management practices, reduction of asset redundancies, risk mitigation, disaster recovery capabilities, etc. Although some of these may be difficult to quantify, the business effects of not addressing them adequately can sometimes be modeled or estimated.

Conclusion

Significant savings can be harvested by leasing services on faster refresh cycles

As one can tell from the HDD failure rates alone, volume-based servers are meant to be 'used hard and retired early'. Technology advances in hardware management and energy management are creating a compelling financial story for adoption of the vendors' latest and greatest. Although it would be difficult (if not impossible) to calculate the exact financial savings of moving to a tighter refresh, rough-cut modeling of the tiers (by year and by type) of the server layers would argue that significant decreases in replacement labor, replacement equipment, energy costs, software licenses, and even server spend could be harvested under such a program.

Market competition between the major server vendors (based on the rapid innovation cycles by the chip manufacturers), creates a future scenario of ever increasing improvements. The studies of 2009 which demonstrated the cost savings of moving to that year's new products were repeated in 2010 with its year of products. The same has started to appear in 2011. The expectation of that market is that each generation

NOTES

"Performance per Watt: Dempsey and Woodcrest"
(David Kanter);
<http://www.realworldtech.com/page.cfm?ArticleID=RW110706025321&p=1>

SPECpower_ssj2008 results:
http://www.spec.org/power_ssj2008/results/power_ssj2008.html

Server Performance Summary – Intel Xeon Processor
<http://www.intel.com/performance/server/xeon/summary.htm>

"Updated: Seagate 1TB Drives Biting the Dust", Gavin Stacy, Jan 15/2009:
<http://www.tomshardware.com/news/seagate-7200.11-failing.6844.html>

SPECjbb2005 page: <http://www.spec.org/jbb2005/>

"Failure Trends in a Large Disk Drive Population", Pinheiro/Weber/Barroso (Google), Feb 2007: http://labs.google.com/papers/disk_failures.pdf

"Disk failures in the real world: What does an MTTF of 1,000,000 hours mean to you?", Bianca Schroeder/Garth Gibson (Carnegie Mellon University), presented at the 5th USENIX conference (ca. 2007); <http://www.usenix.org/events/fast07/tech/schroeder/schroeder.html/index.html>

"Vendor disk failure rates: Myth or metric?", Mary Brandel, Computerworld, April 4 2008; <http://www.computerworld.com/action/article.do?command=viewArticleBasic&articleId=9073158>

"Are Disks the Dominant Contributor for Storage Failures? A Comprehensive Study of Storage System Failure Characteristics"; Jiang/Hu/Zhou/Kanevsky (Univ. of Ill): http://www.usenix.org/event/fast08/tech/full_papers/jiang/jiang.html/index.html

TPC-C benchmarks: <http://www.tpc.org/tpcc/>

For more information, contact:

North America

US	+1 248 253 9000
Canada	+1 866 606 1429

Europe

Ireland	+353 1 238 3451
UK	+44 0 20 3037 5585
Central Europe	+42 0 602 449 467

Australia/NZ

Australia	+61 2 8232 7769
New Zealand	+64 9 357 5876

Asia

China	+66 10 6518 8938
Hong Kong	+852 392 21888
India	+91 22 4230 1200
Japan	+813 3512 7272
Korea	+82 2 3705 8500
Malaysia	+60 3 2059 8833
Singapore	+65 6601 0888