CLABBY ANALYTICS

Research Brief

Does System z Offer Better TCO than Scale-up, Distributed Unix and Linux Servers?

Executive Summary

IBM's System z group is facing two major challenges:

1. First, they've got to convince information technology (IT) buyers that the age old "mainframe" is modern, relevant, and technologically superior to other scale-up servers (particularly high-end Unix and Linux-based systems); and,
2. Second, they need to overcome IT buyer objections to mainframe pricing.

Convincing IT buyers that the venerable mainframe is modern, relevant, and superior to other systems architectures is relatively straightforward. IBM need only focus on mainframe virtualization, security, and its ability to run thousands of Linux virtual machines to do so (*Clabby Analytics* provides a mainframe technology critique in a *Research Brief* entitled "IBM's System z: Technology Update" available at www.valleyviewventures.com as well as at www.clabbyanalytics.com that discusses mainframe architectural advances).

But dealing with pricing objections is a bit more challenging. To get IT buyers to better understand the value that they can derive by purchasing a mainframe, IBM has got to change IT buyer focus away from total-cost-of-acquisition — and refocus buyer attention on total-cost-of-ownership (TCO).

To win the TCO battle, IBM is strongly focusing on demonstrating how its System z mainframe lowers network equipment costs, reduces real estate requirements, consumes less energy than a distributed system environment, requires fewer administrators, and provides more computing capability through higher overall utilization (see Chart 1).

Chart 1 — Potential for Major Operational Cost Savings Using System z Technology

¼ network equipment costs	virtual versus physical
1/25th floor space	400 sq. ft. versus 10,000 sq. ft
1/20 energy requirement	$32/day versus $600/day
1/5 the administration	< 5 people versus > 25 people
Highest average resource utilization	> 70% versus < 15%

Source: IBM — June, 2006

Are these claims valid? In this *Research Brief*, *Clabby Analytics* takes a closer look at IBM's TCO claims. We explore the logic behind these claims; and we provide our opinion on the validity of each claim. And we conclude that:

- IBM's claim to lower network equipment costs by 75% is true because most of System z's communications take place across its internal backplane, and therefore fewer physical switches, hubs and routers are needed to support communications to external processor resources;
- IBM's assertion that it requires less floor space than a group of distributed servers is also true (though perhaps 1/25th the floor space is a bit "aggressive");
- IBM's power/cooling claim also make good sense considering that a System z is a scaled-up system environment that usually runs at greater than 70% utilization (often in the 90%+ range) – whereas a comparably-powered distributed system environment often averages less than 30% utilization while consuming vast amounts of power during idle cycles;
- IBM's assertion to lower administrative/management costs is geographically variable; and,
- IBM's utilization claim is not only valid, it is a major differentiator for the System z architecture.

A Closer Look at IBM's TCO Claims

In this section we take a closer look at IBM's TCO claims in the areas of network equipment savings, real estate reduction, power and cooling, management, and system utilization.

IBM's Network Equipment Savings Claim

IBM claims that it its System z may utilize only 25% of the network components that distributed systems use is a viable claim. The logic behind this claim goes like this:

- Distributed systems need to have one or several Ethernet controllers available to conduct I/O (input/output) between other clustered or connected distributed servers;
- Distributed systems also require more hubs, routers, switches, and bridges than a self-contained, scaled-up System z.
- All of these components accrue additional deployment (cabling), management and maintenance expenses ― and also burn additional power.

IBM's System z exploits its tightly knit communications subsystem and uses its virtual I/O facilities to lower networking and power consumption costs. External hubs, switches, and routers are not required to share data amongst processors. Cost are thusly lower for communications in mainframe environments. And IBM's 75% savings estimate is, accordingly, fully credible.

IBM's Floor Space Savings Claim

Real estate in some markets can cost over $100 per square foot — making floor space an important consideration when calculating TCO. IBM's argument here is that multiple distributed systems have multiple, often large footprints — and hence use-up much more floor space than a self-contained System z. This situation is further acerbated when distributed towers are under-utilized — requiring even more towers to deliver as much computing power as a mainframe (to be discussed later in the utilization subsection).

IBM's argument that System z can reduce TCO because it uses a fraction of the floor space required by multiple, distributed systems appears to be fully credible. But given that distributed servers vary greatly in size and design, there's no "rule-of-thumb" for exactly how much a System z can reduce floor space. IBM's claim to reduce floor space to 1/25 of the area of a distributed system environment is possible – but probably not the norm.

IBM's Power Consumption/Cooling Savings Claim

Over the past year energy costs have fluctuated wildly (so IBM's claim of $32 per day to operate a mainframe versus $600 for an equivalent distributed systems environment is a fluid comparison). Still, IBM's point is well considered. IBM's System z uses highly-efficient power power supplies; does not have to use an external network to move information; and because System z's frequently operate in the 90% utilization range (as opposed to the 20-30% utilization rate frequently found in distributed systems environments), fewer systems are required to do an equivalent processing job.

Because IBM's System z does not rely on external networks for interprocessor communications; and because its design has been optimized for energy efficiency at high utilization rates — IBM's claim of 1/20th the energy consumption rate of distributed system architectures is fully plausible.

Cooling also represents a major cost in the data center. A recent Information Week article (Power Surge — February 27, 2006) claims that enterprises paid about 20% more for electricity last year than they did in 2004 (according to research fime IDC), with rates jumping more than 40% in some parts of the country. Rackspace, a company that provides hosted computing services from five data centers, saw its utility bill increase 65% in 2005. Also noteworthy, space that is leased for $12 to $20 per square foot can cost $60 per square foot to cool.

IBM's mainframe design utilizes fewer components than are required to run multiple distributed systems — and by using fewer components, cooling costs are lowered.

IBM's Management/Administrative Savings Claim

IBM makes a big deal about the high-cost of management and administrative labor as a major contributor to operational costs. To hear IBM tell it, labor is now the highest cost element in distributed systems environments — and administrative staff costs increase in proportion to the number of servers an enterprise deploys. Hence, running a single mainframe should *obviously* cost less from a management perspective (perhaps, as IBM

claims, it may take five mainframe people to handle the workload that it takes twenty-five administrators to handle across a distributed system environment).

We find this claim to be geographically variable. For instance, in Middle East, Indian, and Far East markets, labor costs for administrators are far less expensive than in North American and European markets. Hence, human labor may not be as huge a cost as IBM makes it seem in some markets. Still, in North American and European markets, IBM's claim does have some merit.

It must also be remembered that mainframe administrators often command much higher salaries than administrators with experience on distributed Windows, Unix and Linux systems — hence it may take fewer administrators to manage a mainframe environment, but those administrators may command salaries that are 1.5 to 2X the salary range of Windows, Unix and Linux administrators.

Another point worth considering related to management/administrative costs is in the area of the cost of "human error". Last year's Infonetics study on the cost of downtime in North American Vertical Markets claimed that human error can account for 27% of downtime. And depending on the industry being served, this downtime can run into millions of dollars in lost opportunity and lost labor costs.

One way to overcome human error costs is to automate systems management processes. And IBM's System z does an excellent job of automating systems management (some would argue that the System z is the best in the industry in this respect). By eliminating the need for human interaction with systems, the potential for human introduced errors is diminished.

IBM's Utilization Savings Claim

By making better use of existing computing resources, IBM can argue that it can get more computing done less expensively than on under-utilized distributed systems. To make this argument, IBM claims a 70% utilization average on its mainframes (which we expect is on the low side) — and uses a 15% average utilization rate on distributed servers for comparison (which we also expect is on the low side).

The core argument that IBM puts forward with its utilization rate claim is that it requires only one mainframe to do the work of five, ten, or fifteen large, scaled-up distributed systems. Hence, acquisition costs can be significantly lowered using a mainframe as opposed to having to purchase and manage multiple distributed systems.

IBM also argues that perhaps "average utilization" should not be the comparison point between distributed servers and mainframes. Instead, perhaps the comparison point should be based on "composite peak" performance characteristics. By comparing peak workloads over a fifteen to twenty minute interval on a System z to peak workload characteristics across distributed systems, IBM contends that the above mentioned 15% average utilization rate for distributed servers may actually be lower! Here's why:

Every server in a distributed network does not peak simultaneously. Accordingly, when evaluating distributed system utilization, the sum of the interval data needs to be measured over time (not just certain snapshots). IBM claims that by measuring interval data over time, most enterprises will find that the composite peak of underutilized servers, back-up servers, fail-over servers, QA/test environment servers, and the like will be VERY low (even lower than the 15% utilization figure in most cases) — even when production servers are operated at typically managed peaks of 40-70% utilization.

IBM adds to this mainframe efficiency argument by looking at what happens when additional workloads are added to a mainframe. Adding new workloads in distributed environments often requires additional servers and software (leading to increased acquisition cost which ultimately increases TCO). Adding new workloads to a mainframe is often just an exercise in scheduling.

IBM's case that higher utilization can radically lower computing costs as compared with distributed systems architectures is wholly plausible. Better use of computing resources not only reduces system acquisition costs (by requiring fewer systems to do the same amount of work – as well as through software volume discount purchasing advantages) – it also helps lower management/administrative (operational) costs. Further, IBM can demonstrate that it costs less to add a new application to an existing mainframe than it costs to deploy a new application on a new distributed server.

One Other Point to Consider: Reliability/Availability

One point often overlooked when conducting a TCO evaluation is the cost of downtime. Failures drive-up administrative costs as administrators chase down the root cause for a given failure and ensure that the resource is repaired and brought back on-line.

All high availability systems carry the risk of a long outage as "lock states" occur when a failure is detected — and systems need to be rolled back to that state (which can be time consuming). And loss of processing time often equates to loss of revenue (whether that be loss of customer transaction revenue or loss of revenue due to lost employee work time). For some customers, loss of revenue due to system downtime can reach millions of dollars per hour.

Reliability/availability considerations are very important when considering system TCO. IT buyers need to consider that most of today's systems hardware is pretty reliable. But a closer look at IBM's z/OS operating environment shows a level of software availability matched by no other commercial systems maker. For instance, z/OS functional recovery routines can refresh software modules that may have hung or crashed, rather than requiring a reboot (as is the case in Windows and Unix environments). This kind of software reliability gives IBM's System z a unique and differentiated position when it comes to overall systems reliability – differentiation that should be considered when building a TCO case.

Summay Observations

IBM's System z group is under new management. This group of managers understands very well that it needs to change the impression of the mainframe from that of an aging dinosaur to that of a technology leader. They also know that they need to get IT buyers to turn their focus away from acquisition costs and focus instead on TCO.

To demonstrate their TCO advantages, IBM is focusing strongly on network equipment savings, real estate, power/cooling, management costs, and utilization rates. In each case, they can readily show cost savings as compared with distributed systems architectures. In some cases, these savings are tremendous. In other cases, these savings are variable (for instance, management costs vary widely throughout the world, so operational cost savings related to management costs will also vary widely).

To stay competitive in the long run, we have little doubt that IBM will reduce hardware and software acquisition costs over time. For instance, earlier this year IBM introduced a "baby mainframe", a System z that is priced near $100,000. The company has also worked with SAP to discount SAP software in certain configurations. But don't expect widespread, across-the-board cost cutting. Instead, expect IBM to continue to focus on reinforcing the value that the System z can deliver ─ and expect IBM to strongly emphasize its ***TCO advantages*** as a way to level-the-playing-field between System z and its competitors

Clabby Analytics
http://www.clabbyanalytics.com
Telephone: (207) 846-0498