R+2
THE POWER IS IN THE NUMBERS

Maximizing Hardware Reliability and Minimizing Maintenance Cost

Each year large IT organizations acquire millions of dollars of new hardware. Because competing hardware tends to look alike, contains components that perform the same functions, and are generally in the same price range, they are typically assumed to be equally reliable. This is a flawed assumption. In the fifty year history of the industry it is unlikely that any two competing products have ever been equally reliable or even close to equally reliable. Each time a corporation acquires new hardware, that hardware will alter the reliability profile of its corporate computing system. Put simply, the choice that you might be making today will determine how much hardware failure you will experience over the next four to eight years emanating from that new hardware. As an example, in an IT organization that currently has 1,000 servers; the choice of new hardware could add or subtract 1,000 hardware failures over the next four years simply by choosing the most or least reliable hardware. This is true whether that hardware is virtualized, consolidated, or running one image.

If there is hardware that is significantly more reliable than its competition (and there almost always is) wouldn't it be better if every IT organization knew that? And wouldn't the fact that this knowledge is available to all major hardware buyers force the hardware vendors with less reliable hardware to upgrade that reliability? Wouldn't the vendor with the most reliable hardware today work to improve that reliability to stay in front of its competition knowing that reliability is now measured and knowing that those measurements were now available to all major IT buyers? If I were managing a large IT organization today, I would understand that the best way to maximize my hardware stability is by buying the most reliable hardware available today and by promoting greater hardware reliability in tomorrow's hardware. And I would encourage the vendors not to wait for their next hardware to make my existing hardware more reliable. I would let the vendors of my hardware know that the reliability of that hardware is being monitored, measured and shared with other major IT organizations every month.

When I was the CIO of the fourth largest bank in the US, I asked IBM how I could go about measuring the reliability of the hardware in my data centers. I also asked them if they could provide me with industry standards against which to compare my measurements. IBM responded a week later by telling me that they had no recommendation as to how I could measure the reliability of my hardware and could provide no industry standards on reliability against which my measurements could be compared. On my own, I decided that I could measure the reliability of my hardware by computing the "job failure rate" of each model of each vendor's hardware that I had. And I decided that I could get industry standards by encouraging senior IT managers of other large IT organizations to run my measurement system in their data centers and share their measurements in a common industry pool from which we would all be able to see which vendor's hardware was the most reliable and which were not. While this sounds like something all IT organizations could support today, this experience I am describing occurred forty years ago. I approached three other senior IT executives in the greater New York area, Irwin Sitkin at Aetna, Dave Blackwell at Mass Mutual and Bill Olishar at Prudential. Each agreed to participate, and the first reliability measuring system in the IT industry was born. RPI was soon formed and our pool of contributing IT organizations expanded from 4 organizations to more than 1,000. Gradually that measurement system inspired the vendors to build hardware that virtually ended job failures in the industry.

Some vendors were unable to upgrade their hardware, and those vendors left the industry (Telex, CDC, and Memorex). This user-based measurement system proved to be a highly effective proactive solution to a very real operational IT problem.

What we at RPI are doing now is creating a similar measurement process, only this time instead of measuring the "job failure rate" of mainframe related hardware, we are measuring the "availability rate" and "maintenance rate" of all IT hardware. The "availability rate" is the number of months between disruptive hardware events. A disruptive hardware event is a failure of a component that brings down the hardware or a maintenance action that causes the hardware to be taken down in order to replace that failed component. To demonstrate how wide the spectrum of availability is today, if we look at all processors other than PCs, from mainframes to distributed x86 servers, the availability measurement varies from 9,120 months to 50 months between failures. Occasionally someone will point out that most IT organizations only use their hardware for an average of 60 months, so why should they care that the most reliable hardware has an availability rate of 9,129 months? What those people are missing is that in a pool of 1,000 servers, the hardware has a collective use of 12,000 months per year. And in that situation, a 9,129 availability rate would indicate that that hardware will average slightly more one disruptive failure per year (12,000 divided by 9,129 = 1.3) for the most reliable hardware and 240 disruptive failures (12,000 divided by 50 = 240) for the least reliable hardware. To place that 240 failure number in perspective, that is roughly 10 times more than the number of hardware disruptive failures of all of your mainframe processors, all your RAID storage devices and all your centralized telecommunications hardware in one year combined.

My first priority in developing R+2 was to maximize the stability of the corporate computing system by maximizing the reliability of all IT hardware in that system. That is the reason for computing the availability rate. My second priority was to ensure that the cost of maintenance reflects the substantially reduced need for maintenance, even for the least reliable hardware in use today. In addition to computing the availability rate, R+2 computes the "maintenance rate". The maintenance rate is the number of months between maintenance actions, whether that maintenance action causes a disruptive event or not. It is the maintenance rate that is used to minimize the cost of maintenance. It is time that IT organizations start to reap the significant economic benefits of hardware that is not only very reliable today, but will be substantially more reliable tomorrow if large IT organizations begin to measure it and share those measurements in an industry pool. The maintenance formula that all maintenance providers use today probably made sense fifty years ago when it was first used and hardware failed every month, but today even the least reliable hardware only fails once every fifty months. So why are we still permitting maintenance providers to use that fifty year old formula?

Measuring and proactively managing hardware reliability will bring significant benefit to maximizing hardware stability and minimizing hardware maintenance costs. ***Those benefits can be achieved with almost no effort or expense on the part of your IT organizations.***

Data collection in each IT organization that might have required a substantial amount of work searching through internal data files, instead, under our measurement service, requires no work because the data can be obtained from the maintenance providers. The maintenance providers' data is uniform across all IT organizations (which user

internal data is not) and they will provide your data for all IT hardware they maintain for you each month at no cost.

Once your IT organization has the monthly maintenance providers' data in hand, that data can then be forwarded to RPI. We will combine your data with the data being collected from other large IT organizations across the industry to establish the industry reliability database just as we did for thirty years. Each month RPI will deliver to each contributing IT organization the "Comparative Availability and Maintenance Rate Report". This monthly report serves four primary purposes. IT enables each IT organization to select the most reliable hardware. It encourages all vendors to improve their hardware's reliability, not just in their next hardware, but in their existing hardware. It informs each contributing organization's management how its installed hardware compares with the industry averages. And lastly, that monthly report enables IT management to negotiate a maintenance fee that reflects the increased reliability of the hardware.

Nothing stated above should be interpreted as suggesting that developing the capability to survive hardware failure isn't important. Backup data, backup hardware, and the ability to smoothly and quickly move workload from one resource to another is an essential part of professionally managing hardware. But the ability to survive failure should not be used as an excuse for not taking reasonable and prudent steps to prevent failures in the first instance by promoting greater reliability and by selecting the most reliable hardware. Every disruptive hardware failure impacts the stability of the corporate computing system, regardless of what survival capabilities an IT organization employs. Just because a two engine plane can fly and land on one engine is no reason to buy less reliable engines.

One very effective way to reduce the impact of disruptive hardware failures on the corporate computer system is to reduce hardware failures. And the most effective way to reduce hardware failures is to acquire the most reliable hardware today and encourage the vendors to build more reliable hardware tomorrow.

Jim White
CEO RPI