



The “ticking time bomb” of Virtualization

The “ticking time bomb” of virtualization refers to the problem created when the number of images on a single physical device is increased without a comparable increase in the stability of that device. The key to avoiding this problem is to select the server platform that has the highest degree of stability which means the server platform with the highest availability rate.

Consider the following:

A 12 Month Management View of Stability of the Four Major Hardware Groups

	Centralized Processors	Distributed Servers	Telecommunication Hardware	On-line Storage
Disruptive Events *	0	1,818	1	2
Availability Rate**	9,120 months	66 months	1,540 months	1,600 months

Notes

*A disruptive event is one where a hardware failure causes the hardware to stop performing or where the maintenance action requires the hardware to be taken down.

**The Availability Rate is the numbers of months between disruptive events and is used to define the stability of competing architecture.

What the table discloses

Of the four hardware management groups, one will generate 99+% of the hardware disruptive events that interfere with the stability of the corporate computing system and that is the group that is currently undergoing virtualization.

The huge disparity between the availability rate of distributed server hardware that causes 99+% of the hardware disruptions of the corporate computing system and the centralized hardware (centralized processors, on-line storage, telecommunication devices) that causes less than 1%, is that centralized hardware is designed to continue operating when it experiences hardware failures or when failing components are being replaced while distributed hardware is not. It is a matter of the philosophy of the architecture of large hardware where disruptions are prevented versus the architecture of small hardware where they are not.

Centralized hardware that produces less than 1% of the corporate computing system’s disruptions uniquely shares five important characteristics that distributed hardware does not. It is these characteristics that give centralized hardware such high availability rates:

- Extensive use of component redundancy.
- Failed components that can be replaced without interfering with the operation of the hardware.
- Hardware monitoring by “Call-Home” systems that communicate problems directly to trained vendor engineers.
- Engineering Changes (ECs) that are provided over the life of the hardware that promote stability and extend the useful life of the hardware.
- Quick maintenance response is almost never required.

A Typical Comparison of Distributed Servers versus a Centralized Server

Function	IBM	HP	Dell	Sun	Fujitsu	IBM system z processors
Server months	46,489	40,096	6,267	29,817	2,855	9,129
# Maint Actions	321	605	61	468	44	19
Maint. Rate	144.82 months	66.27 months	102.74 months	63.71 months	64.88 months	480.00 months
Disruptive events	305	575	58	445	42	1
Availability Rate	152.42 months	69.73 months	108.05 months	67.00 months	67.97 months	9129.00 months

The above table shows a sample of the availability rates (the number of months between disruptive events) of five x86 based processors that comprise a portion of the processors in the distributed world of hardware and a mainframe which is one of the processors that comprise the centralized world of hardware. The full R+2 server processors report, of which the above table is just a small sample, shows the availability rates of all server options. The full report is divided into two parts, distributed and centralized processors. Distributed processors includes all x86 based platforms, and centralized processors include, but are not be limited to, large Unix-based systems, Superdome, HP's Non-stop systems and IBM mainframes. The above sample server report shows a stability range from one disruption every 67 months to one disruption every 9,129 months (760 years).

Using the availability rate to select the most stable hardware

Using the shortened version of the report shown above, if an IT organization had 10,000 servers, the difference between using IBM's distributed servers and HP's distributed servers would be the difference between having 787 disruption events per year or 1,721. And the difference between using a distributed HP server and a mainframe would be the difference between having 1,721 disruptive events per year and 13. In the coming years all IT organizations are going to be faced with the problem of which server platform architecture to virtualize onto. There may be good reasons for staying in the distributed world and equal reasons for moving to the centralized world, but for the IT management that is committed to maximize the stability of the corporate computer system, regardless of which of those two architecture they choose, there can be no reason for not choosing the most stable hardware within that architecture.

The key to maximizing the stability of the corporate computing system is to acquire the most stable hardware. And the key to acquiring the most stable hardware is to have access to the availability rates of all IT hardware.

Quoting from Joe Clabby's paper referring to IT organizations that choose hardware without knowing its stability rating: *"The big strategic mistake that your company is making in systems management and maintenance is that you're basing your decisions on a void — a dearth of management and maintenance statistics. And because you don't have those numbers, your company may be choosing the wrong server architecture, the wrong server platform, and the wrong server vendor."*

[For more information on this subject be sure to read "The Major Strategic Mistake Your Enterprise is Making in Systems Management and Maintenance" under Clabby Analytics Papers on this site.]