



Linux System Shipments and Virtualization Less Can be More!

By Dan Kusnetzky, Principal Analyst

O V E R V I E W

Kusnetzky Group clients often question what is the real size of the market for Linux? They're hoping to map out their IT budgets for the coming years without accidentally betting on the wrong platform. They're confused by all the different metrics they see presented in press releases and presentations offered by industry research firms. One view makes the market look like its expanding while another; equally reputable, report might paint a bleak picture of a declining market.

While I'm pleased that the confusion often leads decision makers to my door, I must admit that the real answer is likely to be complex. It's clear to me that a bit of explanation is in order of both types of market measurements and what each type really means.

W H A T A N A N A L Y S T S E E S D E P E N D S U P O N W H A T I S B E I N G E X A M I N E D

The size and trends in a complex market, such as the market for open source software in general and Linux system software in specific, can be challenging to uncover. Analysts at research firms such as Gartner Group, International Data Corporation and others publish hundreds of reports annually in the attempt to present a clear, accurate and unbiased view of market dynamics.

The challenge, of course, is that all of these firms are looking at different things, in different ways, and, of course, come up with different statistics and forecasts. It's not that one firm is right and the others are wrong, it's that there are many views of any market and *all of them are correct*. What's a decision maker to do? A good first step is to understand what is being counted, what is being excluded and how is the count being done.

I'm reminded of a few quotes.

"Statistics: The only science that enables different experts using the same figures to draw different conclusions." Evan Esar (1899 - 1995), Esar's Comic Dictionary

"A lie told often enough becomes the truth." Lenin (1870 - 1924)

"A lie can travel halfway around the world while the truth is putting on its shoes." Mark Twain (1835 - 1910), (attributed)

"There are three kinds of lies: lies, damned lies, and statistics." Benjamin Disraeli (1804 - 1881)

S L I C I N G , D I C I N G A N D M A R K E T S T A T I S T I C S

Analysts who are attempting to understand a complex market, such as the market for Linux software usually try to follow the hardware or follow the software. Furthermore, they may examine shipments or the revenue produced by

Document #20070606 The Kusnetzky Group © 2007

The Kusnetzky Group is an independent supplier of marketing services to suppliers and user organizations and suppliers in the systems, virtualization and open source technology markets. Suppliers of hardware, software and virtualization technology are among those organizations. The opinions presented in this document are based upon our research, our personal experiences and actual use of technology regardless of whether this document or the supporting research were sponsored by one or more of the Kusnetzky Group's clients. This document may not be copied in whole or in part without the written permission of the Kusnetzky Group.

those shipments. Top analysts will do all of these things and compare the results. Let's look at all the different things that might need to be counted.

Hardware Shipments — How many systems were shipped worldwide? This is often segmented into client and server shipments

Software Shipments — How many copies of software were shipped worldwide? This is often segmented into Client Software and Server Software. Top analysts also segment each category into paid and unpaid shipments.

Hardware Revenues — How much revenue did the shipments of hardware produce for the suppliers? This is often segmented into client shipments, server shipments and mainframe shipments.

Software Revenues — How much revenue did the shipments of software produce? This is often segmented into revenues produced by client software and server software.

If I'm counting correctly, a decision maker might have to purchase a minimum of 9 different reports to get a partial answer to their question. What they don't often know is that even when taken together, these nine reports wouldn't fully explain the Linux market dynamics. Here are some issues that are seldom addressed in most research reports.

Since open source software, such as Linux, can be replicated as many times as desired by an organization, one downloaded copy of Linux might turn into a deployment of 10s, 100s, or, perhaps, 1000s of systems. They might be client systems or server systems.

A new application may not lead to the purchase of a new system. An older system might get a "field promotion" to Linux from another operating environment.

Virtual system deployments might rise dramatically while physical system shipments slow or decline.

As systems become more powerful, an organization may choose to purchase a small number of powerful systems rather than a larger number of less powerful systems. How can an analyst measure what work is being done by these systems? The larger machines may actually be performing more work for the organization than the more numerous smaller machines.

Recent industry reports have pointed out that virtualization has caused Linux system shipments to slow or even decline. Before those in the Linux industry start trembling and thinking that the future is bleak, the key questions they should considering is whether trend is something good? Is it something bad? Or is it just one of many metrics decision-makers should consider?

WHAT IS VIRTUALIZATION?

Virtualization has become a major catch phrase being bandied about in the business and IT-oriented media today. Usually, the media mistakenly equates virtual machine technology with virtualization. Furthermore, it is often presented as something new and different rather than a well-established group of technologies with a long track record of success in data centers around the world. What is it really? Virtualization is the use of hardware and software technology to present a logical view of resources to developers, administrators, operators and system users. This logical view is often strikingly different than the actual physical view.

What does this really mean? System users may see a single physical computer as if it were many different systems running different operating systems and application software. Or they may be presented the view that a group of systems are actually a single computing resource. Other virtualization technology may allow individuals to access computing solutions with devices that didn't exist

when developers created an application. It may also present the image that long obsolete devices are available for use in the virtual environment even though none are actually installed.

WHAT IMPACT WOULD ONE EXPECT OF THE USE OF VIRTUALIZATION TECHNOLOGY?

The market impact one would expect to see depends upon what type of virtualization technology is being considered. Although I'm not going to present an exhaustive analysis that encompasses all six layers of virtualization technology the Kusnetzky Group examines, a review of the impact of one layer, processing virtualization software, will shed some interesting light on the subject.

Processing virtualization uses hardware and software to break the link between an application, application component, system service or whole stack of software and the underlying machine(s). Let's look at each type of software in this category and project what increasing use of that type of software would have on the market for Linux software.

- ☒ **High performance/Grid computing software** — software that allows applications or data to be segmented (sometimes the word "decomposed" is used) into units that can be run simultaneously on many machines to reduce the amount of time necessary to execute applications. There are several different communities using this type of software. Some call it "grid computing," some call it "parallel processing," and others make references to the marching brooms in the Disney film "Fantasia." This one-application-to-many-machines approach is typically found supporting scientific, research, geophysical, weather and content creation applications. If this software saw increased use, the number of Linux blades or Linux systems would increase dramatically. These types of configurations often deploy hundreds or, perhaps, thousands of independent systems. Since processing power, not scalability is the focus of this type of computing, these systems would be small but, powerful.
- ☒ **High availability/clustering software** — software that is used to monitor an entire application stack or service. If the application or service slows down or stops completely, a replica of the application or service is started on another machine. This approach is commonly used to support a monolithic function such as a collaborative application or a database management engine. If this type of software saw increasing use, the number of Linux blades or Linux systems would increase modestly. Since scalability and reliability are the goals when this type of processing virtualization software is deployed, the systems would, in all likelihood, be large.
- ☒ **Virtual machine software** — software that gets beneath and then encapsulates an entire stack of software from the operating system, data management software, middleware and applications themselves. This software then allows multiple "capsules" or virtual machines, each running its own operating system, to run on the same piece of physical hardware. This approach is commonly deployed as part of a server consolidation strategy. Interestingly enough, it might also be used to support client consolidation. As this type of virtual processing software saw increasing use, the number of Linux blades and Linux systems would decline even though the amount of work being done by these systems would increase. The goals of these deployments are optimization of system usage. Fewer, but much larger, systems would be deployed.

- ☒ **Partitioned operating systems** — this software segments a single operating system so multiple stacks of software can share the same operating system without interfering with one another or even knowing that other stacks or Zones exist. As with virtual machine software, this type of software is often part of a server consolidation strategy. The impact on the market for Linux would parallel the impact caused by virtual machine software.

If we added up the impact of the increasing use of all of these types of software and considered the fact that Linux has moved to a mainstream operating system used by organizations of all sizes, in all markets and in all geographic areas, we would soon come to understand that the following things.

- Fewer but, much larger, systems would be seen if the metric was system shipments.
- If the metric was value of systems rather than shipments, the value would tend to rise.
- If the metric was Linux shipments, the result would be that the number of instances of Linux in use would rise (if non-revenue copies plus those producing revenue were included in the count).
- If the metric was revenue produced by those Linux software shipments was the focus of the analysis, the result would be mixed. Virtualized environments lead organizations to deploy more no-cost replicas rather than many more commercial copies.

WHAT DOES THE FUTURE HOLD LINUX?

It is clear that virtualization technology is going to play an ever-increasing role in the data center and in systems close to individuals. The increasing need for reliability, agility, security and management are directing organizations to consider and then deploy virtualization.

It is also clear that Linux has moved from being a tool found largely in scientific and academic circles to being a mainstream operating environment. This move has changed how Linux is deployed from being whatever system is the cheapest to what machine is large enough for the task at hand.

Suppliers of Linux technology are likely to see larger and more complex installations but, fewer physical systems.

Published research from IDC, Garter and others bares this out.