IDEAS

actionable intelligence
www.ideasinternational.com

Are product lifecycles really shrinking? How long should users keep servers in production?

See page 2

TECH TRENDS MONTHLY

March 2010 Newsletter

REGULARS

Ask IDEAS

IDEAS Insights

FEATURE

How Old Are Your Servers?

Using product lifecycle information from Competitive Profiles, IDEAS analyzed the product lifecycles of nearly 2,500 different server products going back almost 20 years to see if the perception of ever-diminishing product lifecycles matches the reality.

Tech Trends Monthly is free!

To subscribe click below link or visit:
www.ideasinternational.com/misc/msc060.aspx

subscribe

How Old Are Your Servers?

Gary Burgess, SVP, Research & Operations

If one were to ask IT users what challenges they face in managing their IT infrastructure, the chances are that managing the rate of change, and/or keeping up with technology advancements, would be high on their lists. Increasingly, users are finding themselves on a continual upgrade treadmill, just to keep up with the latest technology. Every time they purchase a new piece of technology – either hardware or software – it seems to be quickly superseded by another release.

Ideas International (IDEAS) has been monitoring the server and storage infrastructure space for many years. Recently, IDEAS released a new feature in the Competitive Profiles tool that provides deeper insight into the lifecycles of various product families. Using this product lifecycle information from Competitive Profiles, IDEAS analyzed the product lifecycles of nearly 2,500 different server products going back almost 20 years to see if the perception of ever-diminishing product lifecycles matches the reality.

Are product lifecycles really shrinking? How long should users keep servers in production? How often should they upgrade to keep in step with technology advances that may help secure that competitive edge for their business?

Two-Year Server Lifecycles

For our analysis, the product "lifecycle" for a server begins when the server is announced, and it ends when the server is officially withdrawn from marketing. Using this definition, our analysis shows that since 1991, the average server market lifecycle is around 27 months. The data for this analysis included servers classed as entry (up to two-socket), midrange (up to eight-socket), and high-end (above eight-socket). Both x86 and UNIX servers were included, but mainframes were not. Blades were also included in the analysis sample, when they came onto the scene around 2001.

Chart 1 (next page) provides more detail, looking at the annual lifecycle averages since 1991 by the four product categories described above (Entry, Midrange, High-End, and Blade). The horizontal axis represents the year a product was announced and the average lifecycle of all servers announced for each product type in that year is plotted, generating a line for each product type across the featured period.

The data does not include current products, which by definition have not yet reached the end of their marketing life. This is the main reason why the chart only goes up to 2007. Beyond that, the sample sizes of obsolete products are too small to be statistically relevant.

The outcome for any single year is not important, given the potential for minor statistical anomalies along the way. An example is the high-end server product type in 1997. The peak is caused by the Fujitsu Trimetra range of servers being withdrawn from marketing during 2004 and 2005, with lifecycles in the range of 87 to 98 months each. Taking these Trimetra products out of the calculations brings the average lifecycle of the remaining products down to 36 months, which is more in line with this product type during this period.

The trends over time are more interesting. Generally speaking, the Blade, Entry, and Midrange servers have had very similar lifecycles, averaging between 24 to 27 months each. The High-End servers have usually had slightly longer lifecycles – around 34 months on average. However, the High-End server product type lifecycles have been on the decline more recently. From 2005 onwards there is a degree of convergence, with all product types having similar lifecycles.

(Continued on page 3)

Are product lifecycles really shrinking? How long should users keep servers in production? How often should they upgrade to keep in step with technology advances that may help secure that competitive edge for their business?

Chart 1 – Server Lifecycles

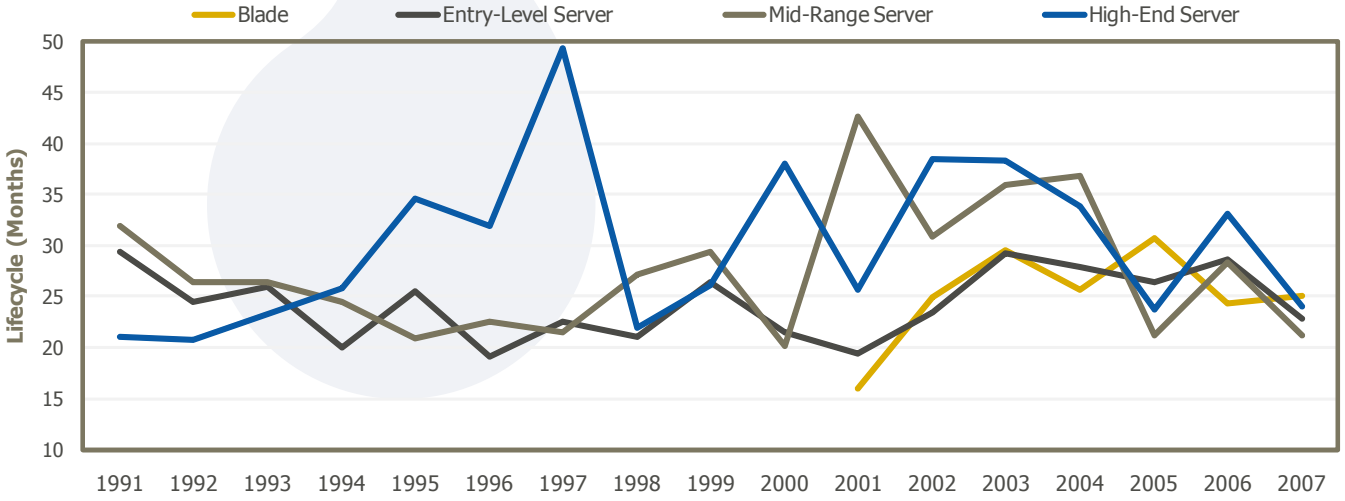

Source: Competitve Profiles

The Effect of Commoditization

IDEAS speculates that commoditization is the main reason for this convergence of product lifecycles. As product margins have declined, the need to control costs has risen. So, in addition to other strategies, developing products that share common components across the performance spectrum is a logical outcome and that is what has been happening.

In the past, higher-end products were generally very different from lower-end products in terms of their design and the types of components they used. But today, high-end and low-end products are increasingly sharing common components. As new technology is released, there is a desire to flow that technology through to all products as soon as possible. Hence, the growing number of shared components results in a natural alignment of product lifecycles across the different product types.

The Hidden Factor

To some extent, the overall outcomes appear to show that product lifecycles have not in fact been shrinking, at least over the period from 1991 to 2007.

However, there may be a hidden factor that is not reflected in the data: Within the past 10 years or so, products have tended to see multiple processor generations during their lifecycle, whereas prior to that time, a new faster processor usually came packaged as a brand new server model in the family. Hence, much of the perception of change can probably be attributed to these processor bumps during the life of a server, while the rate of change of the actual server – the container housing the processors, I/O, and memory, etc. – has had a relatively constant lifecycle across all that time.

However, it is important to note that although processors can be upgraded, most of the other server technology is fixed for the product life, such as I/O support and memory technology. With the focus on Green, or Efficient, IT in recent years, power consumption and heat dissipation are becoming as critical as processing power. Hence, the requirement to move to new boxes that exploit more efficient technologies will take on even greater relevance.

The IDEAS Bottom Line

In review, the average marketing life of a server is just over two years. IDEAS has received anecdotal feedback that the average lifecycle in the field of an x86 server is just short of three years. That number seems to correlate with what this research has revealed about the sales lifetime. As such, IDEAS offers the following takeaways:

- Users should take these lifecycles into account when assessing servers to purchase. If the server you have been reviewing has been in the market for over 18 months, the chances are it is about to be replaced by something better.
- Managing a server for a three-year working life seems to be a good fit with vendor server refresh cycles.
- A three-year lifecycle may also make sense for higher-end servers given the greater commoditization and the trend toward similar product lifecycles between high-end and lower-end servers.
- The three-year lifecycle also correlates well with the common three-year point of sale service offerings. Vendors will typically offer an effective discount for the up-front three-year hardware maintenance service commitment, compared to purchasing via an annual contract during the server lifecycle. ■

Ask IDEAS

Question: What kinds of applications are most suitable for cloud computing?

Answer: While no industry-wide agreement has yet emerged on the definition of cloud computing, there is consensus that the benefits of cloud computing are easiest to achieve with scale-out applications, in which workloads are deconstructed into multiple tasks that can run independently on clusters of multiple computers. In their purest form, cloud-based applications assume that individual tasks can fail, or fail to complete, so they implement ways to allow the overall service to continue processing, even if individual nodes of the cluster fail. By handling the availability of workloads in software, many reliability features are no longer needed in hardware for operations to continue processing in the event of component failure. With fewer reliability features in individual compute nodes, the cost of cloud server hardware can be pushed much lower than that of traditional systems, which are typically designed for enterprise computing. This design point in cloud-optimized hardware is emerging as a key component of the economic benefits promised by cloud computing.

Scaling out plays a key role in many aspects of cloud computing. For example, as users deploy private cloud infrastructure, they typically employ virtualization to control resources, consolidating virtualized server, storage, and network resources into pools of virtual infrastructure that can be accessed by workloads on demand. With such virtual infrastructure, key operational aspects such as scalability, reliability, and even power consumption are typically handled by dynamically shifting virtualized workloads around between multiple hosts. Applications based on Infrastructure-as-a-Service (IaaS) and Platform-as-a-Service (PaaS) often depend on virtual machines, which means they will scale out by taking advantage of the distributed workload management underlying the virtual infrastructure.

Scaling out can also be applied to other usage models that are associated with cloud computing. For example, web serving is perhaps the most common service offered by hosting providers, many of whom are aggressively pursuing cloud computing approaches. Web hosting is relatively straightforward to scale across multiple servers through the use of network load balancers. Software-as-a-Service (SaaS), which is also commonly associated with cloud computing, usually depends on some variant of web hosting, and thus benefits from scaling out. Finally, many HPC workloads have been migrated to clusters of servers and standardized interconnects, aided by operating system and HPC application awareness (as evidenced by the presence of clusters in the TOP500 list of the world's fastest supercomputers).

While some third-party cloud computing services do support platforms that scale up, such as POWER hardware offered with IBM's Computing on Demand service, most of the leading cloud platforms and services are clearly optimized for scaling out. Therefore, users who are interested in pursuing cloud computing should focus on those workloads that have been proven to work well in scale-out environments. ■

Do you have a question for IDEAS?
AskIDEAS@ideasinternational.com

. . . users who are interested in pursuing cloud computing should focus on those workloads that have been proven to work well in scale-out environments.

IDEAS Insights

Blog Bites (http://www.ideasint.blogs.com)

IBM is one of the few **vendors of x86 servers** to perform its own **engineering** of the major system **components.** While most other x86 server **suppliers use Intel's standard chip sets** as the building blocks for complete server systems, **IBM has differentiated** its x86 servers with **unique technology** since acquiring Sequent in the late 1990s.

IBM eX5 Servers Pump Up Memory and Cores in Pursuit of Virtualized Workloads
Tony Iams | March 2, 2010

read more»

Now, **Fujitsu** has introduced the **PRIMERGY CX1000,** an x86 server platform designed specifically to host **cloud computing workloads.** The system is **optimized for scale-out computing** with hundreds to thousands of nodes, simultaneously targeting a **minimal space footprint, innovative cooling, and simplified operations** with a single integrated approach. A highlight of CX1000 is its new **cooling mechanism,** called Cool-Central.

Fujitsu PRIMERGY CX1000 Brings Turbine to Power Cloud Computing
Tony Iams | March 17, 2010

read more»

About Ideas International

Ideas International (IDEAS) is the leading global supplier of comparative intelligence on enterprise IT infrastructure. Servicing both IT vendors and large-scale IT users, IDEAS products focus on: servers; storage; virtualization and consolidation; and software and infrastructure services. IDEAS online products are backed up by expert industry analysts who can also deliver specialized advisory and consulting services.

For more information about Ideas International, please visit:

www.ideasinternational.com

IDEAS INTERNATIONAL

Americas

Ideas International, Inc.

800 Westchester Avenue
Suite N337
Rye Brook, NY 10573-1354
USA

Tel +1 914 937 4302
Fax +1 914 937 2485

Asia/Pacific and Worldwide Headquarters

Ideas International Limited

Level 3
20 George Street
Hornsby, NSW 2077
Australia

Tel +61 2 9472 7777
Fax +61 2 9472 7788

Europe, Middle East, Africa

Ideas International Europe

The Courtyard
Lombard Street, Abingdon
Oxon, OX14 5BJ
United Kingdom

Tel +44 (0) 1235 462 890
Fax +44 (0) 1235 462 891