‘Hyperversity’ Driven by Technical & Cost Differences

GABRIEL CONSULTING GROUP

Our recent annual x86 Data Center Surveys have shown that the average x86 enterprise customer is using three different hypervisors to virtualize their workloads. Our latest survey shows that most customers aren’t simply test-driving them in hopes of finding the “One True Hypervisor.” They have solid reasons, primarily focused on cost and technical differences, for using multiple solutions and will continue to do so for the foreseeable future…

Our recent report “*Hyperversity Rages On*” examined how real-world data centers are using x86 hypervisors. We found that the average respondent’s data center currently runs three different hypervisors. A sizeable number, close to 20%, have four or more mechanisms; only a third use just one virtualization package. We’ve invented the term ‘hyperversity’ (a combo of ‘hypervisor’ and ‘diversity’) to give us a quick and catchy way to refer to it.

The most widely used hypervisor is, not surprisingly, from VMware; it was present in a little more than 80% of the data centers we surveyed. Microsoft’s Hyper-V was in second place, in use in 40% of data centers. KVM and the Citrix Xen variant were each being used by about a third of our respondents.

All of this data, incidentally, comes from the 2011-’12 edition of our annual x86 Data Center Survey. Close to 350 enterprise data center personnel answered a wide range of questions about how their organizations are coping with business/technology challenges, how they’re responding to various trends, and how they regard the major x86 vendors. We’ve included some demographic data from the survey at the end of this report.

In this edition of the survey, we dug a bit deeper into this multi-hypervisor question. Are customers just kicking the tires on various solutions? Testing them to see which fits best in their infrastructures, with an eye toward standardizing on one package? Or do they plan to use multiple hypervisors forever? (In the IT industry, that’s probably about five years).

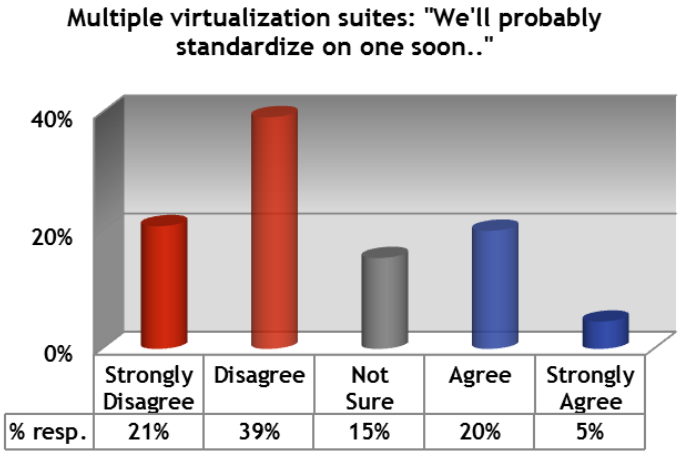

It’s pretty clear that the vast majority of our respondents aren’t just kicking the tires on multiple hypervisors. Only a quarter say that they expect to settle on a single hypervisor in the near future.

© 2002-2012 Gabriel Consulting Group, Inc. All Rights Reserved

Fully 60% explicitly say they *won't* be standardizing anytime soon – with 21% saying they strongly disagree, in essence a "Hell no" response. A smallish 15% aren't sure whether they'll standardize or not.

So the vast majority aren't standardizing now. Why not? Let's discuss a few reasons why a customer wouldn't standardize.

One of the major differences among hypervisors today is acquisition cost and total cost of ownership. VMware, with its vSphere functionality and vCenter management suite, is firmly anchored at the top of the market in terms of both pricing and functionality.

After a painful revamp of its pricing, VMware now charges customers based on both sockets and memory. Licenses range from $1,500 to $3,500 per processor on the ala carte plan – plus annual support subscriptions from $500 to $1,250.

Microsoft's Hyper-V is at the other end of the scale – it's included with the Server 2008 operating system. Citrix Xen and KVM are both open source products available as either free versions or commercial editions supported by major vendors (Citrix and Red Hat, respectively).

Pricing for commercially supported editions of these packages typically comes in well below VMware price levels – and not by coincidence. KVM, for example, is fully licensed and supported by Red Hat for around $500 (standard) to $700 (premium) per processor per year.

Cost factors into every business and data center choice, of course. And the costs can mount quickly even in a small or mid-sized environment when you begin to add up the per-socket and annual support costs. When it's all totaled, the sticker shock can be profound for the unprepared.

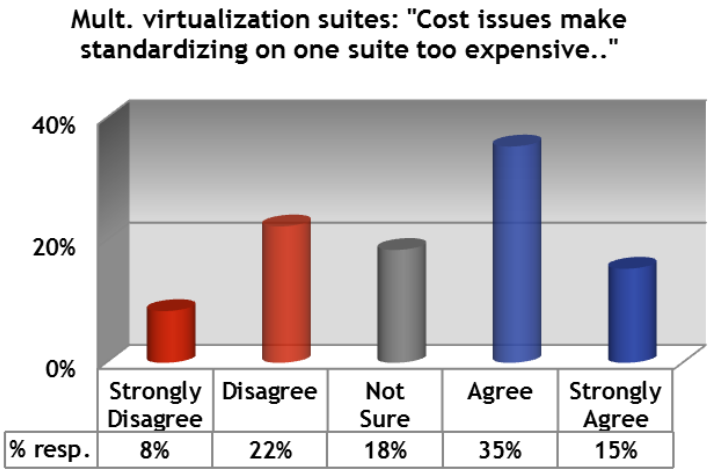

In our survey, half of our respondents cited cost as a major reason for not standardizing on a single hypervisor anytime soon. Just under a third don't see cost as something that's keeping them from standardizing, while 18% aren't sure.

Cost is certainly an issue when it comes to hypervisor standardization. However, most of our respondents say that technical factors are what's keeping them from using a single hypervisor solution.

© 2002-2012 Gabriel Consulting Group, Inc. All Rights Reserved

Technical Issues Drive Hypervisor Diversity

At a high level, there isn't a massive amount of difference between hypervisors when it comes to capabilities and functions. However, there are some technical differences that can make one particular hypervisor a better choice than the others for a specific situation.

The most obvious example of this is the match-up between Microsoft's Hyper-V solution and Windows Server systems and applications. Hyper-V is included with the Windows Server o/s and is fully supported by Microsoft, so it's natural for Microsoft-centric data centers to use it for their Windows applications. But VMware, Xen, and KVM are also fully capable of supporting Windows guest VMs and many customers choose to virtualize with these packages rather than Hyper-V.

On the Linux side, KVM has a bit of an advantage over VMware and Xen. KVM is the only hypervisor that is integrated into the standard Linux kernel. This gives KVM an advantage in that it inherits native Linux scalability and device support, plus it uses native Linux mechanisms to manage memory, scheduling, and the like. Since it's part of the Linux o/s, KVM virtualization doesn't have to be integrated with the customer application stack – it just works right out of the box. This also explains why some customers find that it's a bit easier to deploy and manage.

Xen and VMware are 'bare metal' hypervisors – meaning that they run on a special software layer that sits under the host operating system(s). Bare metal hypervisors need to use their own unique mechanisms to portion out system resources and handle scheduling tasks for their guest operating systems. Both solutions are managed using their own unique APIs, while KVM can be managed with standard Linux commands. KVM can also be run as a 'bare metal' standalone hypervisor if desired.

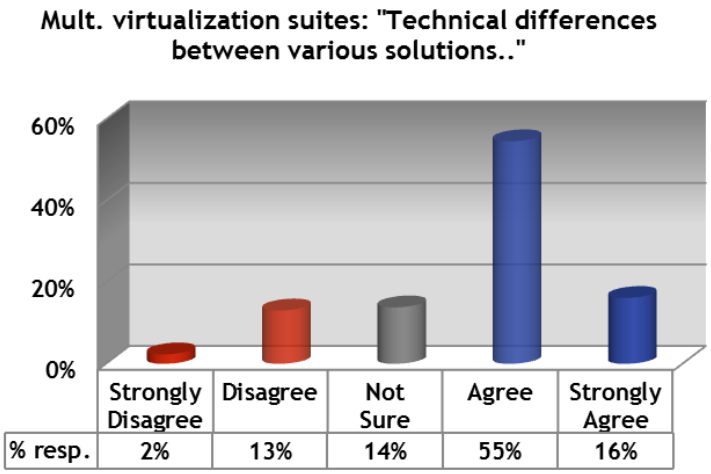

Technical issues and differences are the biggest factors driving hypervisor diversity. As can be seen on the chart at left, more than 70% of our respondents cite technical differences as the major reason for running multiple hypervisor solutions.

Only a small number of respondents, 15%, say that technical differences don't make much – well – difference to them.

© 2002-2012 Gabriel Consulting Group, Inc. All Rights Reserved

Private cloud computing is one area where the technical differences between hypervisors may be coming into play. In parsing our survey data, we isolated heavy users of each of the major hypervisors to see whether any significant differences were revealed. We defined 'heavy users' as respondents who said they were using a particular hypervisor on 'many systems' or had standardized on that solution.

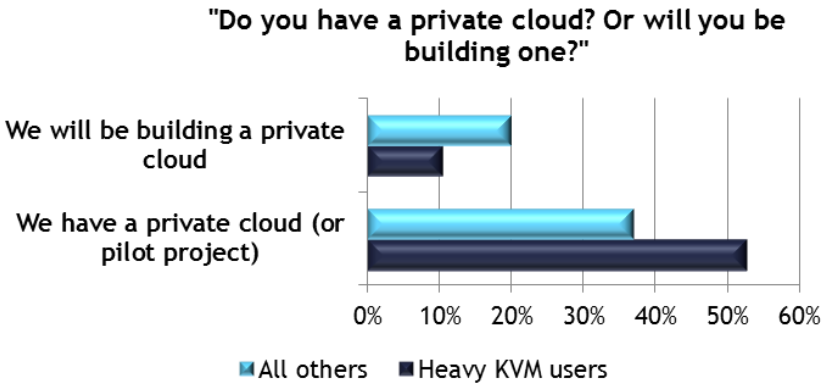

One item that jumped out at us was the difference in private cloud adoption rates between heavy KVM users and those who are primarily using other hypervisors.

More than half of the KVM-centric organizations (53%) already have private clouds or functioning pilot projects in place vs. 37% of the non-KVM shops.

While this is only a snapshot in time, it's an interesting data point to consider, and it raises some questions. Is there something about KVM that makes it more ameneable to driving private clouds? Or is there something about the folks who gravitate toward KVM that makes them more likely to move to private cloud computing more quickly? We can't draw solid conclusions from the data we have, but it's something that bears further observation and research.

Does 'Not Soon' Mean Never?

While 60% of our enterprise respondents say that they don't plan to standardize on a single hypervisor in the near future (first chart in this report), does this mean that they *never* see themselves doing it?

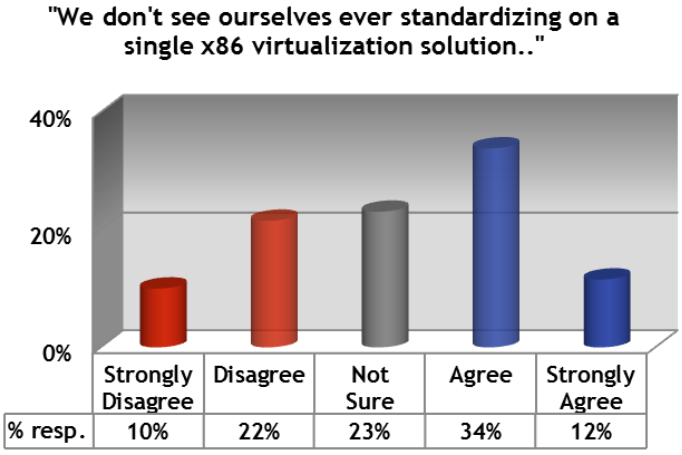

Almost half, 46%, say that they won't ever standardize on a single hypervisor solution. Almost a third of them believe that they will eventually land on the "One True Hypervisor," but this is only marginally greater than the 25% who (in the first chart) said they were going to settle on a single provider soon.

What we take from this is that the number of customers who are still kicking hypervisor tires is pretty small. Most have decided that they're always going to run multiple solutions.

© 2002-2012 Gabriel Consulting Group, Inc. All Rights Reserved

Summary

While some believe that hypervisors will soon become commodities – and thus not all that interesting – we disagree. Virtualization solutions and the hypervisors they rest upon are key control points on a system and in the enterprise infrastructure. Hypervisors aren't interchangeable, and standardizing on a single hypervisor will result in either financial and/or technical trade-offs.

According to our research, single hypervisor users tend to be smaller organizations with fairly homogeneous IT infrastructures. Larger organizations with more complex IT requirements are using different hypervisors in order to get the "right" technical solution for their needs.

While half of our respondents say that cost differences are a major factor in their choice to run multiple hypervisors, a much larger 71% say that technical issues are a bigger factor.

Since these technical issues aren't going to go away anytime soon, we expect to see the cold war between the hypervisor vendors continue to bubble away for the foreseeable future.

This document may not be reproduced or transmitted in any form by any means without prior written permission from the publisher. All trademarks and registered trademarks of the products and corporations mentioned are the property of the respective holders. The information contained in this publication has been obtained from sources believed to be reliable. Gabriel Consulting Group does not warrant the completeness, accuracy, or adequacy of this report and bears no liability for errors, omissions, inadequacies, or interpretations of the information contained herein. Opinions reflect the judgment of Gabriel Consulting Group at the time of publication and are subject to change without notice.

© 2002-2012 Gabriel Consulting Group, Inc. All Rights Reserved

GCG 2011-2012 x86 Data Center Survey Demographics

Our surveys are aimed at actual data center workers. We believe that they have the best handle on how their organizations use technology in general as well as specific hardware, software, and services. They have multi-product and multi-vendor experience, and they understand what their businesses and IT operations are planning for the future.

Organization Size - Total Employees

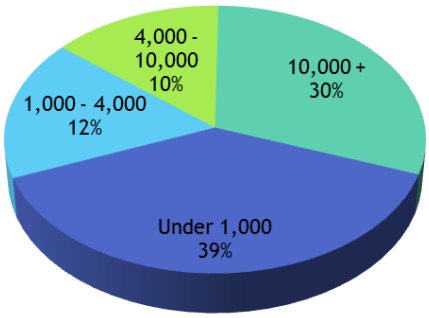

This survey was in the field from the last half of 4Q2011 into the second quarter of 2012. Respondents came from GCG's survey data base of previous participants and from targeted advertising. The total number of respondents to this survey was 345, with 40% of respondents in mid-size and large organizations of 4,000 employees and above.

Geographic Location

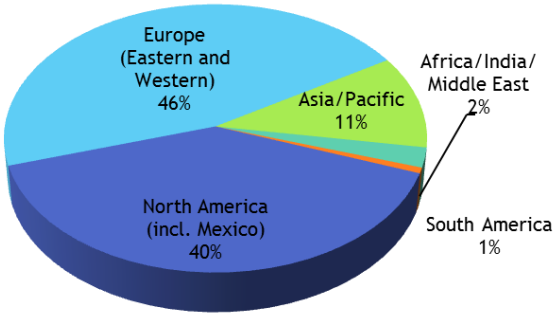

This was a global survey; 46% of respondents hailed from Europe, 40% from North America, and 11% from Asia/Pacific.

Annual IT Budget

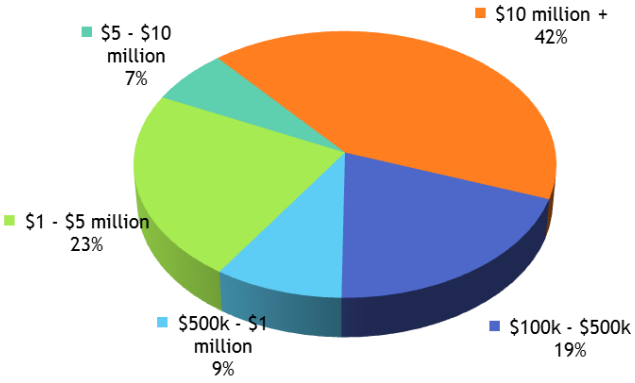

As part of our demographics, we asked respondents for their organizational annual IT spending. This included hardware, software, and services provided by third parties. 42% reported spending of more than $10 million, and more than half spend at least $5 million per year on IT products/services.

Respondent Responsibilities
(multiple selections allowed)

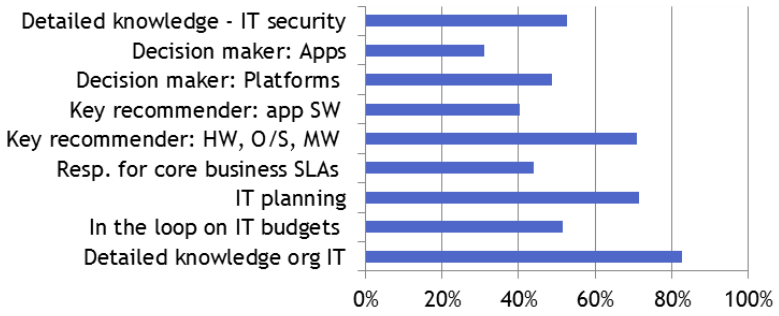

We also asked each respondent to describe their personal knowledge of their organization's IT infrastructure. As can be seen on the chart, our respondents are key players in their data centers and bring high levels of knowledge and experience to the survey.

© 2002-2012 Gabriel Consulting Group, Inc. All Rights Reserved