Server Virtualization in Context

Rationalizing and optimizing your x86 server estate

Dale Vile and Jon Collins, Freeform Dynamics Ltd, September 2009

One of the hottest topics in the IT industry at the moment is virtualization, particularly in relation to the x86 servers that have so often proliferated to unmanageable levels. While the principle of server consolidation based on the latest virtualization technologies is now accepted, how far have organizations progressed in this area? And based on adoption experiences, what are the practical considerations when dealing with server proliferation?

KEY FINDINGS

Wake-up call: organizations generally have more physical servers than applications

Feedback from a recent study suggests that the IT infrastructure in larger organizations is often supporting several hundred, if not a thousand or more, applications, with even smaller businesses supporting software portfolios in the 10 to 50 application range. While this may be familiar, the wake-up call is that applications are generally outnumbered by the physical servers on which they run. As a result, 85% of respondents highlight existing or emerging issues with server proliferation.

Server proliferation is a function of cultural as well as technical factors

Historically, new applications have been installed on their own dedicated hardware, regardless of whether the full capacity of a server is required – this avoids conflict with other applications, and enables each box to be tuned to run an application in an optimum manner. However, the dedicated server approach reinforces the (administrative and political) expectation of business stakeholders owning everything associated with the applications they fund, with the server and other equipment allocated to their own cost centre.

The consequences of server proliferation are real, but can be tackled

Server sprawl has a direct, negative impact on routine activities such as patch management, application provisioning, and general monitoring and management of performance. This has a knock-on effect with regard to operational overheads and associated costs. Server proliferation also goes hand in hand with poor server utilization and power and space related challenges, which not only translate to elevated costs, but can also constrain development and growth. Those who have server proliferation under control demonstrably suffer significantly fewer problems in all of these areas.

Virtualization technologies are key to driving improvements

Quantitative and anecdotal evidence suggests that there are clear and tangible benefits to be gained from the implementation of virtualization technology to consolidate and rationalize x86 server estates, and experience in the mainstream is being accumulated rapidly. With the solution landscape still developing, however, it is important to monitor the way in which offerings are evolving in terms of pricing, bundling and capability, e.g. something that looked current a year ago might not do so today.

Adoption experiences highlight the importance of forward planning

When adopting any new technology, it is important to ensure that new problems are not being created for the future, e.g. for the unprepared, unwanted proliferation of physical servers can so easily be replaced by virtual server sprawl. Understanding implementation and management best practice, and planning accordingly, will reduce the risks and enhance the returns from your virtualization activity.

The study upon which this report is based was independently designed and executed by Freeform Dynamics and executed in collaboration with The Register news site. Feedback was gathered via an online survey of 301 IT professionals from the UK, USA, and other geographies, and an interactive 'reader workshop'. The study was sponsored by Microsoft.

Introduction

Virtualization is a hot topic, but one full of apparent anomalies. Speak to any mainframe veteran, and they will be quick to point out that the principle has been around for 30 years or more and is extremely well proven. Speak to some vendors, and they will tell you that virtualization is now also well established in the distributed systems environment, especially in relation to x86 servers.

While virtualization might now be far more generally available than in the past, the reaction among mainstream IT professionals has been mixed. The majority are very positive about what virtualization can potentially do for them in terms of rationalizing their x86 Windows and Linux server estates, for example, but even those with experience often say there is still a lot to be learned about putting together investment cases, designing and configuring virtualized environments, and implementing operational best practice.

Against this background, this report summarizes the results of an IT practitioner study completed in September 2009, during which feedback was gathered, predominantly from those with experience of x86 virtualization, via an online workshop and survey (see Appendix for more details) conducted by Freeform Dynamics in association with *The Register* (www.theregister.com). Based on this input, we review the background and primary rationale for the adoption of x86 virtualization technology, then move on to consider some of the implementation practicalities. The aim is not only to assist those new to the area to get up to speed, but also to help practitioners with more experience to checkpoint their progress and approach with reference to insights gathered from their peers.

To kick this discussion off, let's start by looking at one of the most common issues facing IT departments today, server proliferation, which is typically the prompt for virtualization initiatives.

The problem of server proliferation

In the cut and thrust of operating a busy IT department, it can be all too easy to lose sight of how tactical additions to the IT environment can accumulate over the years. While we tend to pay attention to megaprojects and the big investments in core applications such as ERP and CRM, deployments such as the implementation of a more modest software package or the execution of a quick bit of development work can slip under the radar.

Through all of this activity, however, the amount of software accumulated over time is often significant. In larger organizations, for example, it is normal for the IT infrastructure to have to support a software portfolio made up of several hundred, if not a thousand or more, applications, with even relatively small businesses at the other end of the spectrum typically finding themselves in the 10 to 50 application range (Figure 1).

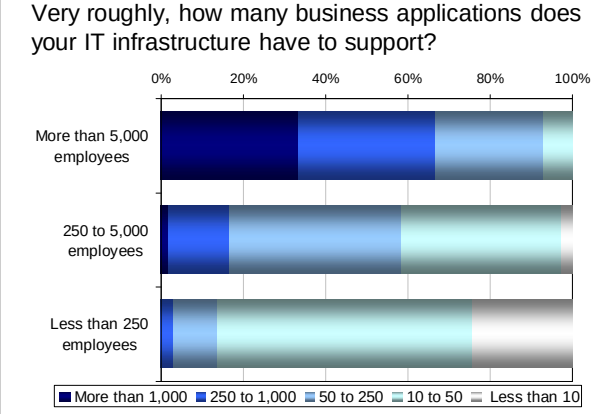

It is all too easy to forget the amount of software that accumulates over time that needs to be supported by the IT infrastructure.

Figure 1

With this picture in mind, it is interesting to consider the amount of server kit required to drive all of this software. Clearly, there will be some big systems in the mix, such as the aforementioned core

business applications, larger externally facing websites, etc, that may need a whole server landscape of their own to support them. At the other extreme, we then have workgroup or departmental applications supporting just a handful of users, but each still needing some horsepower to drive them.

When you net all this out, it is interesting to look across different organization sizes at how the average number of applications (shown above in Figure 1) compares to the average number of servers (Figure 2), and some may be surprised that the latter is actually greater than the former.

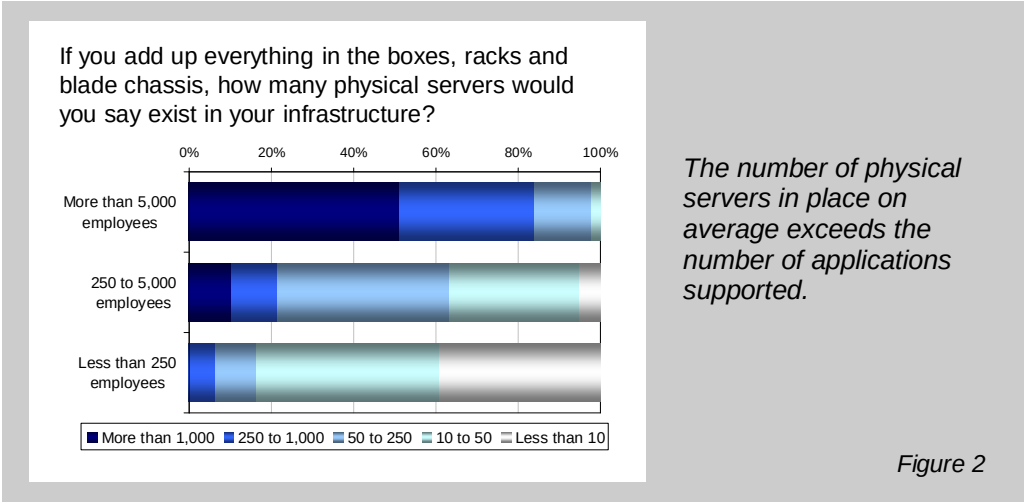

In other words, while the challenge of application proliferation may be great, the challenge of server proliferation is even greater. Given these stats, it is completely understandable that over 85% of those providing feedback in our survey allude to past, ongoing or emerging issues in this area (Figure 3).

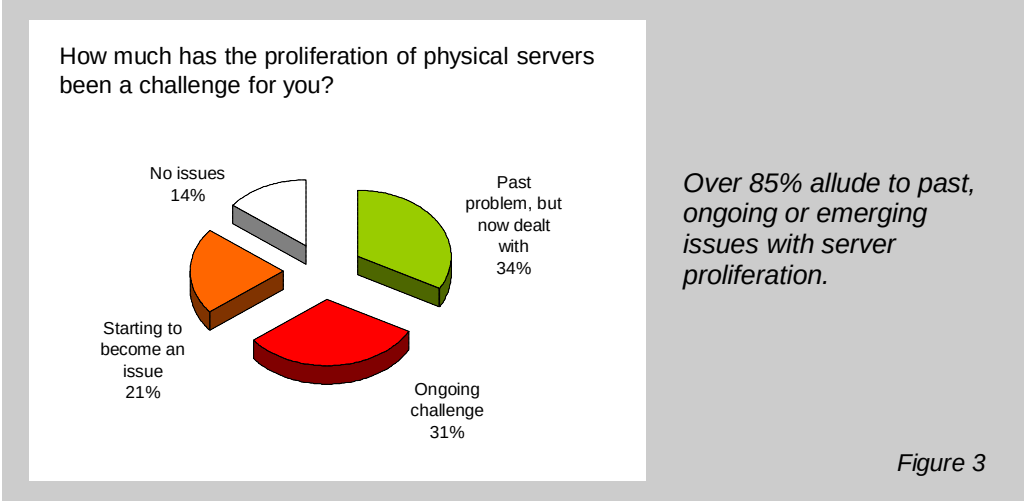

One contributing factor to the level of server proliferation is the fact that beyond the applications themselves, we have all of the horizontal systems capability that needs to be in place, ranging from directory services, through networking services such as DNS (Domain Name Service) and DHCP (Dynamic Host Configuration Protocol), to security and access control, monitoring and management infrastructure, and so on. A degree of server infrastructure is clearly required to support all of these.

We then have the aforementioned 'big' applications that require multiple servers to drive them, and these too can contribute significantly to the amount of server equipment that accumulates.

While the expansion of server estates is understandably and unavoidably driven by such requirements, there is another way in which servers accumulate which is much less desirable.

Causes of unwanted server proliferation

A less desirable contributor to the server proliferation phenomenon is the traditional practice of implementing each application on its own dedicated hardware.

To illustrate the point, consider a situation in which a new application is required to support a small group of users in one part of the business.

In this scenario, once the relevant software has been selected or developed, it needs to be deployed on the appropriate hardware. While its needs from a performance perspective might be very modest, with the expectation of it only 'tickling' a low-end server (even at peak load), we elect to install it on dedicated hardware anyway. One reason for this is so it doesn't conflict with other applications (or vice versa), and another is so the box can be configured and tuned to run the application in an optimum manner.

The dedicated server approach, however, also caters for the expectation of business stakeholders to own everything associated with the solution they have funded. In line with this, the server, and any other equipment required to drive the application is thus allocated to the stakeholder's cost centre, and everything is kept straight from an administrative and political perspective.

In anything but the smallest of organizations, this has been a very familiar routine over the years, to the extent that it just became accepted as the way things were done. With this in mind, it is not surprising that the feedback from our study confirms the kind of behavior and underlying causes we have been discussing (Figure 4).

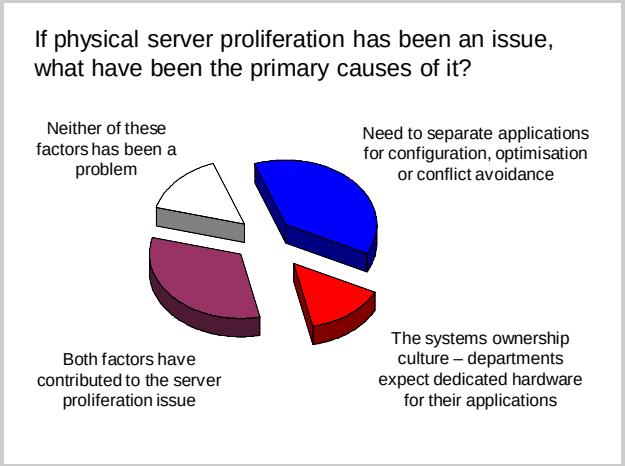

Hardware/software dependencies and the systems ownership culture are both acknowledged as important contributing factors to the problem of server proliferation.

Figure 4

As we can see, there is variation between organizations in the relative contribution of perceived technical dependencies and the influence of cultural and administrative factors, but overall, some 85% of those encountering server proliferation challenges acknowledge one or both issues.

So, while organic growth of the IT infrastructure is partly driven by legitimate needs for increased capacity, the traditional tying of hardware to software for the reasons we have been discussing is primarily responsible for what many describe as the 'server sprawl' that exists within many computer rooms and data centers today. And a key attribute of this sprawl is under-utilization, i.e. a lot of server hardware doing not a lot of work.

The consequences of this, apart from the intuitively negative wasting of capacity, are significant.

Consequences of server sprawl

By way of context, when we look at the operational side of IT, it is easy to identify a broad range of general issues and challenges that are all confirmed by the participants in our study to one degree or another (Figure 5).

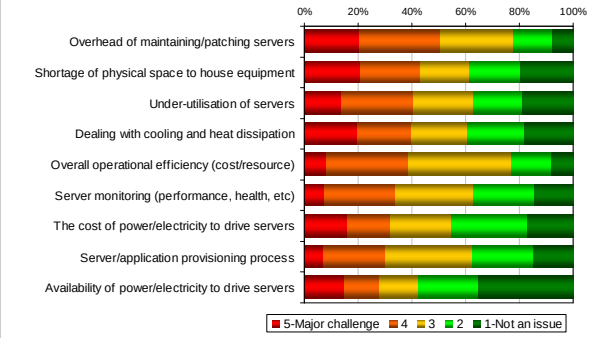

When we look at the operational side of IT, it is easy to identify a range of challenges.

Figure 5

What we are looking at here is based on a rating of the degree to which each area is a challenge on an intensity scale of 1 to 5.

This is interesting enough, but if we home in on the percentage of respondents indicating a more extreme rating of 4 or 5, which (it is reasonable to assume) relates to a higher intensity challenge, we can see some marked differences between those that have server proliferation under control, and those for whom it remains an issue (Figure 6).

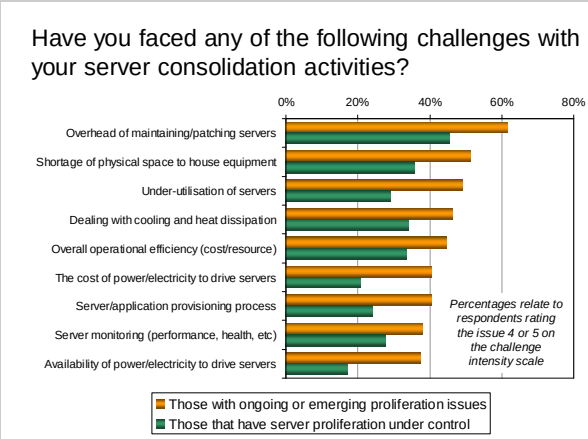

There are some marked differences between those who have proliferation under control, and those who do not.

Figure 6

What's clear from this is that proliferation of server hardware has a direct negative impact on routine operational activities such as patch management, application provisioning, and general monitoring and management of performance.

According to these findings, proliferation also goes hand in hand with the server utilization issues already mentioned, along with the related power and space challenges.

Of course by turning our analysis and reasoning around, we can argue that as most organizations today are suffering or beginning to suffer from server sprawl, the above picture can be interpreted as providing an indication of the kind of areas in which we can drive significant improvements through consolidation and rationalization within our server estates.

To put it another way, this nicely sums up the case for virtualization, which today offers a primary mechanism to achieve control over the server proliferation challenge.

Virtualization adoption and progress

So, what can we learn about how organizations are adopting server virtualization? The majority of those participating in our survey (92%) had implemented x86 server virtualization to one degree or another (Figure 7), so we can gain a great deal from their insights.

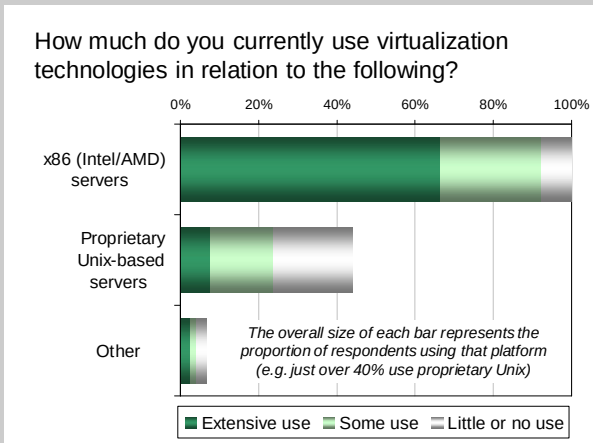

The participants in our study were mostly IT professionals with actual experience of deploying virtualization solutions.

Figure 7

We should at this stage point out that online surveys of this nature tend to exaggerate adoption levels due to the 'self selection' principle, i.e. those with more knowledge of and/or interest in a topic are more likely to participate. Indeed, we know from other research with no inherent bias conducted with the same overall pool of IT professionals (i.e. readers of *The Register*) that adoption levels as of today are more around the 50% level.

In the context of the current study, however, the bias towards those with experience is a significant benefit, as input on practicalities is much more likely to be well informed rather than based on guesswork or supposition.

Having said this, even within our relatively progressive sample there is still a significant subset that are only just beginning their virtualization journey, and when we look at the progress made in terms of x86 server consolidation, it is clear that regardless of the extent of adoption, the job of dealing with the sprawl is still ongoing or yet to be tackled in most cases (Figure 8).

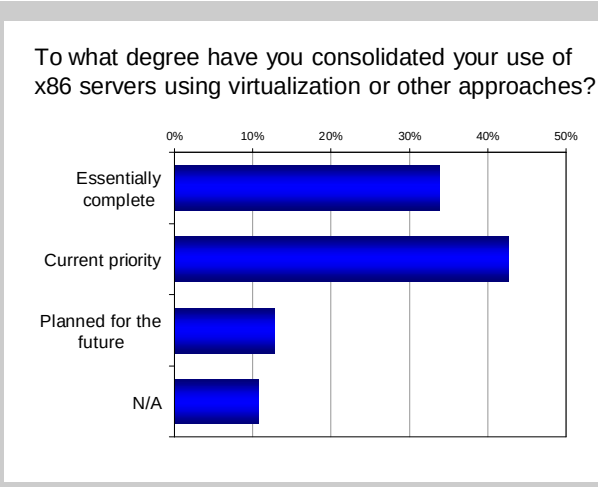

There is still much to be done on server consolidation, even within the progressive group participating in this study.

Figure 8

This 'work in progress' view we are picking up here is consistent with an intention to ramp up the use of virtualization technology in the short to medium term in 70% of cases (Figure 9).

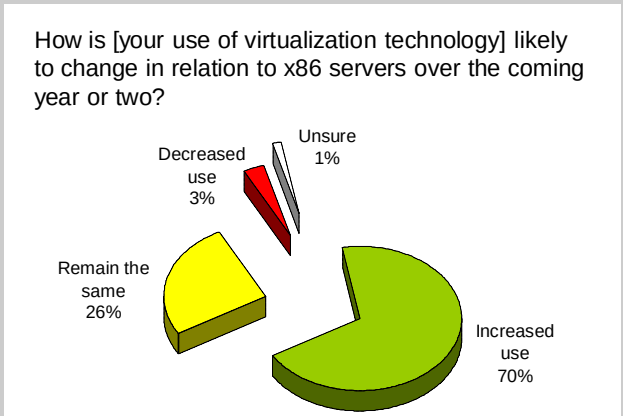

The majority are looking to increase their use of virtualization technology over the short to medium term.

Figure 9

Given this, it is appropriate that we now spend some time looking at implementation practicalities.

Implementation practicalities

While the data we have been considering so far gives us a feel for the overall rationale for server virtualization and the way its use is developing, any experienced IT professional knows that the devil is always in the detail. Let's therefore drill down and consider some of the key practical questions.

Which types of server can be virtualized?

To this point we have talked about server virtualization in a very generic way without really considering which specific types of server are likely candidates for targeting. One way of thinking about this is in terms of the nature of the workload being run - application servers, web servers, database servers, and so on. We can also look at the class of application being supported - core applications, departmental/workgroup solutions, etc. Either way, the general consensus is that pretty much all types of server are up for grabs from a virtualization perspective (Figure 10).

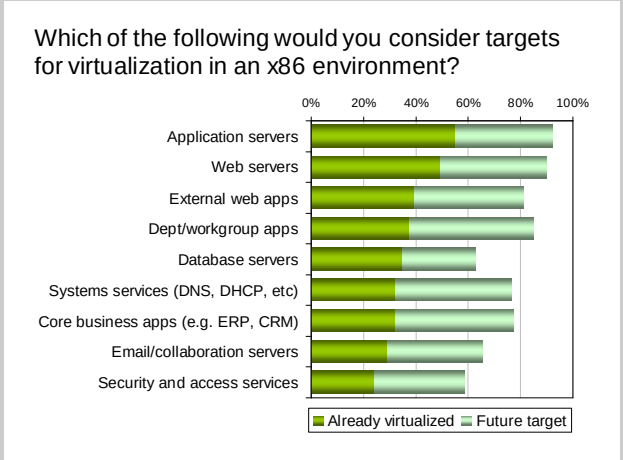

The general consensus is that all types of server are potential candidates for virtualization.

Figure 10

When it comes to targeting, it therefore makes more sense to focus on factors such as utilization rates, support overhead, and so on when considering where virtualization technology can be deployed to best effect. The only caveat is to look out for I/O intensive workloads that may not require a huge amount of processing capacity, but can easily become network or disk bound when co-hosted with other applications on the same physical server. It is not that such workloads cannot be virtualized, as hypervisors can allow effective allocation and prioritization of resources, but it is an issue to be aware of when mixing and matching software on the same physical machine.

What kind of consolidation ratios can be expected?

Regarding consolidation ratios, the claims made by both vendors and practitioners vary considerably. Here, for instance, are some quotes taken from the online workshop run as part of our study:

> *"We're running about 15 VMs per server: a mix of Windows and FreeBSD mostly, some high power (e.g. mail), some low power, but there's still plenty of room for more."*

> *"I work for local government, and we have consolidated close to 25:1 on x86 Windows servers over a 2 year period."*

As we can see, it is not uncommon to hear people talk about ratios of 15:1 and upwards, but when we look across the sample as a whole, consolidation ratios of between 5:1 and 10:1 are more typical (Figure 11).

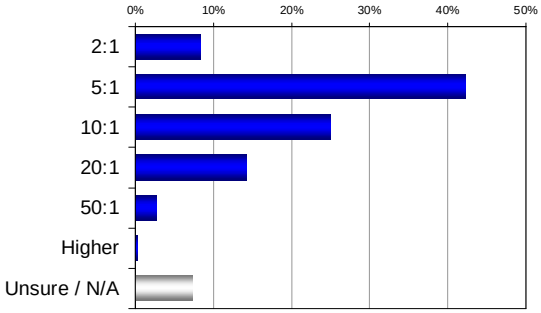

Consolidation ratios of between 5:1 and 10:1 are typically achieved with x86 server consolidation.

Figure 11

Nevertheless, there is little doubt that significant consolidation of hardware estates occurs quite naturally as part of a virtualization initiative. The first and most obvious benefit here is hardware cost saving, as summed up by the comment:

> *"We also have been re-deploying virtualized hardware servers (if decent spec/age) instead of purchasing new hardware servers."*

Fewer servers to support the same application load clearly also translates to a reduction in operation overhead, and a boosting of operational effectiveness.

How does virtualization achieve operational benefits?

During the online workshop, there was some discussion of the difference between running multiple applications on the same operating system instance to improve server utilization, versus hosting them on the same machine, but each in their own partition. This provided a good prompt to flush out some thoughts on exactly how the virtualization approach delivers operational improvements:

> *"[Hosting multiple applications on the same operating system instance] is all well and good until two vendors' packages conflict. Or you have to tell the people using the other 10 applications you installed 'Sorry, rebooting the server, nothing to do with your stuff, it's the other guy's, but it's all on the same box...'. Virtualization minimizes the hardware while still keeping each vendor's tech support happy and minimizing conflicts and single-points-of-failure."*

> *"25 apps on the same OS install, with overlapping ports, libraries, web servers, drivers... one vulnerability on 1 app and a hacker has all of your infrastructure, nice. Need to do a hardware update, just your entire business down while you re-install 25 apps. Have a poor app with a memory leak = crash entire business for a while, instead of 1 app down."*

It was also interesting to hear experienced practitioners articulate how very clear benefits can be achieved from a systems resilience and disaster recovery perspective:

> *"High availability server hardware for EVERY system: Previously we could only afford servers with redundant power supplies or RAID 10 for some systems. Most lived on RAID 5 or 1. With virtualization we have servers with dual power supplies, multiple network cards and RAID10 with battery backup."*
>
> *"The benefit of virtualization to our disaster recovery solution can't be overstated. We backup virtual machine folders on to disk and tape. Simple, fast, no expensive 'backup agents' or other complexity required and can be restored onto any hardware."*
>
> *"The big driver for virtualization was DR. With the previous mix of physical servers meaningful DR was impractical. Now it's easy."*
>
> *"We have lots of operational flexibility. We can test software upgrades to production servers by taking a copy and testing in a separate virtual environment and a snapshot before doing the upgrade in the production environment."*

The big message here is that when putting business cases together to justify the investment of resources and budget into a server virtualization initiative, it is important to look beyond the capital and operational cost savings and consider the positive impact on service levels and risk management.

What should we be prepared for in terms of costs?

While the benefits of server virtualization are very tangible as we have seen, it is clearly going to be necessary to understand the cost side of the equation. This is important as the basic technology can be obtained for little or no incremental investment. The core hypervisor element, for example, can often be simply downloaded for free, or activated as part of an operating system bundle, but this low cost of entry is both a blessing and a curse.

The danger is that consolidation activity can start in a relatively *ad hoc* and informal manner, but quickly escalate because of the benefits we have discussed to the point where the free or bundled solution is outgrown and/or other consequences of virtualizing more of the server estate become apparent. And with the IT department now dependent on the technology and the approach, the unanticipated costs that arise become both a practical and political issue.

Here are some example comments from practitioners who have seen additional costs sneak up on them in this manner:

> *"Didn't spend enough on the disk storage, and now we have run out. Upgrading this is going to cost lots, possibly more than the initial roll out. We [also] maxed out our memory at the time. Sadly we have used pretty much all of it and again this will cost lots to upgrade."*
>
> *"...we may save tons of money on the server hardware, but we spend the savings on the software and supporting [network and storage] hardware... every four physical servers we buy needs to come with a networked disk, and a switch (or two)."*
>
> *"The only blockage at the moment is cost"*

In terms of the need to spend money on software (alluded to in the second comment here), there are two main things to consider.

The first is that the free or bundled solutions sometimes come with restrictions in terms of the hardware configurations and workloads they will support, so it is important to understand the restrictions and the cost of upgrading to the full 'enterprise' version if you think this may be required down the line to meet your needs.

And when investigating this area, don't assume that all vendors take the same approach to packaging and pricing features and functions, or even that information gathered some time ago is still valid. From a supplier perspective, it's a highly competitive market that is just breaking into the mainstream, so expect a lot of movement on the commercial front as the level of activity increases and players jostle for position.

The other consideration from a software licensing perspective is systems management. Even if you have a good management toolset in place for dealing with physical servers, it may or may not be capable of working effectively when things become virtualized. As your use of server virtualization escalates, monitoring virtual machines, moving them around, maintaining them in terms of patching, configuration, etc, is probably not something you will want to continue doing manually. Indeed, while we have previously talked about the problems of physical server proliferation, it is easy to fall into the trap of allowing virtual machine sprawl to occur. The importance of management automation is therefore not to be underestimated.

With this in mind, when embarking on or looking to scale up a virtualization initiative, is important to think ahead and anticipate your management needs. One of the decisions to be made will be whether to extend your existing management infrastructure and tools to include the virtual environment. Management solutions exist, for example, that allow both physical and virtual assets to be monitored and administered from the same console using the same set of policies. This has obvious advantages as it is unlikely you will want to virtualize everything, in which case you will end up with a mixed environment. Having said this, some vendors argue that the needs of the virtual environment are special, and recommend setting up parallel facilities for the management of virtual servers. Which way to go is a hotly debated question in the industry right now, so it is worth talking it through with both existing and new suppliers. In some cases, such conversations may even prompt a refresh or replacement of your existing management systems, which will have even more of an impact on costs.

The message here is to think ahead beyond your initial experimentation or pilot activity, and make sure you fully understand the longer term cost implications of the choices you are making.

What are the considerations from a security perspective?

Another area often overlooked when activity starts in a less formal manner is that of security, but as the following comments from experienced practitioners illustrate, it is an important area that needs to be considered:

> *"While providing an abstracted view of the hardware, an operating system also brings some baggage. One bag contains security holes that the applications that run upon the operating system do not have. The same is true for hypervisors. The more features that are integrated into the hypervisor such as management facilities, the more vulnerabilities arise. An operating system places a bunch of applications at increased risk and a hypervisor places a bunch of guests and even more applications at risk."*

> *"Threat Mapping and Risk Analysis – A broad threat mapping exercise should be undertaken to look at the level of risk and threats associated with virtualization that is specific to the environment / business market that you are in."*

It is beyond the scope of this report to provide detailed recommendations, but during our study we received some interesting pointers on the kind of things to think about:

> *"Consideration will be needed in terms of patching and virus control being in a centralized environment; also what Denial of Service and recovery methods that will be required to manage a virtualized estate."*

> *"How will virtualization impact on industry best practice advice in terms of 'segregation of administrative duties' e.g. virtualization administrators assuming the role of traditional network engineers as we move to more layer 2 devices becoming a virtualized commodity?"*

This last comment is interesting in that it highlights the potential impact of virtualization on roles and responsibilities within the IT department, the message being that it is not just about technology.

How likely is it that pitfalls will be encountered?

When we look at the frequency with which challenges have arisen as a result of consolidation and virtualization activity, we see some of the things we have mentioned already called out at the top of the list, such as the system's ownership culture, the problem of virtual machine sprawl, and the systems management and cost impact of scaling up activity (Figure 12).

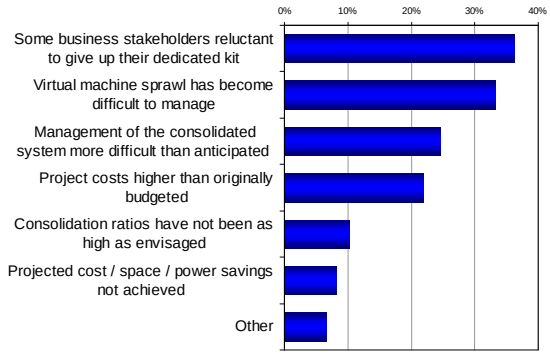

It is important to be aware of potential issues associated with the system's ownership culture, and the operational and cost impact of scaling up virtualization activity.

Figure 12

But let's not get hung up on the notion that these potential 'gotchas' are major blockers. For every person that reported a challenge in each area, two or more said they hadn't experienced the issue. In some cases, this may be because they have not driven far enough down the virtualization road to discover some of the problems, but it is important to acknowledge that the general spirit of the feedback received from practitioners is that the pros far outweigh the cons, however far you push things.

Discussion

Based on the feedback from mainstream IT professionals reported in this document, it is safe to say that x86 server virtualization is now well accepted as playing an important part in the future of IT service delivery. The core benefits in terms of capital and operational cost savings are clear, and other benefits to do with service level enhancement and operational risk management are also becoming better understood.

The link between application proliferation and server proliferation cannot be overstated – but it is worth noting that in this report we have dealt with the challenge from the server perspective. We know from other research that significant benefits can be gained from reviewing the application portfolio and addressing issues such as duplication, fragmentation and under-utilization when it comes to the applications themselves. However, virtualization offers more than a salve – many organizations will have no choice but to keep going with their existing applications, and virtualization can yield benefits whatever the state of the application portfolio.

All the same, the evidence suggests that server virtualization is still a work in progress. It is therefore important to monitor the way in which solutions are evolving from a pricing and bundling as well as a capability perspective. Taking the time to understand how best practice is developing, and planning accordingly, will also enhance the returns from your virtualization activity, both in the shorter and longer term.

With this in mind, we hope this report has been useful, and in the spirit of community under which Freeform Dynamics operates, we would welcome any feedback you might have.

Appendix

Sampling and Methodology

The research upon which this report is based was partially gathered through an online survey executed via a popular news and information site. The sample was made up predominantly of IT professionals with an involvement in virtualization. A good cross section of industries was included and the distribution by organisation size was as follows:

How large is your organisation in terms of number of employees?

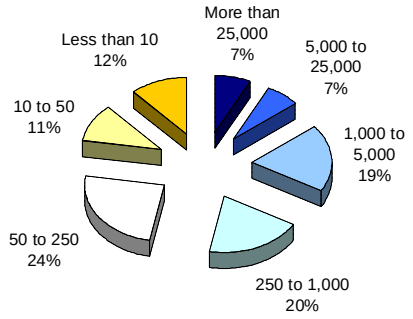

The sample contained a broad cross section of industries, with respondents coming primarily from the USA and UK.

Total sample = 301

Note that the usual caveats to do with online research apply to this study, namely that respondent profiles are self declared and the 'self-selection' sampling process is likely to have skewed the sample towards those with an interest in or knowledge of x86 server virtualization. Neither of these factors, however, can reasonably be expected to have had an impact on the conclusions outlined in this report.

Beyond the survey, input from IT professionals was elicited via an online workshop, in which discussion points were published, then comments invited. The advantage of this approach is that respondents were able to express themselves freely in the context of the discussion, and the quotes included in this report are examples of some of the feedback received.

If you are interested in looking at the discussions that took place in more detail, please visit the workshop summary and navigation page at www.theregister.co.uk/software/virtualization_workshop/.

About Freeform Dynamics

Freeform Dynamics is a research and analysis firm. We track and report on the business impact of developments in the IT and communications sectors.

As part of this, we use an innovative research methodology to gather feedback directly from those involved in IT strategy, planning, procurement and implementation. Our output is therefore grounded in real-world practicality for use by mainstream business and IT professionals.

For further information or to subscribe to the Freeform Dynamics free research service, please visit www.freeformdynamics.com or contact us via info@freeformdynamics.com.

About Microsoft

Founded in 1975, Microsoft (Nasdaq "MSFT") is the worldwide leader in software, services and solutions that help people and businesses realise their full potential.

For more information on Microsoft's virtualization solutions, please visit http://www.microsoft.com/virtualization/.