White Paper on Consolidation Ratios for VDI implementations

Executive Summary

TecDem have produced this white paper on consolidation ratios to back up the return on investment calculations and savings achieved over the past 7 years with server virtualisation projects.

The IBM System x3850 M2 was able to scale to a phenomenal ratio of 227:1 [1]

This can result in a 75% reduction in costs when compared to a single desktop PC [2]

Our experience extends to the original high end servers from IBM and HP and this paper specifically shows what is possible with the high end IBM x3850 M2 and x3950 M2.

These two models have raised the bar for high consolidation ratios above what Dell or HP have to offer and this is demonstrated with the findings and reports in this white paper.

Key findings of the tests and simulations we carried out are:

- High consolidation ratios on large physical servers is more economical that blade (scale out) environments when licensing costs and power consumption are factored into the overall total cost of ownership
- It is possible to calculate a total power consumption number per user and determine how much can be saved or diverted to other IT requirements
- VDI is one of the easiest cost justification solutions available at the moment

[1] TecDem had VMware's permission to exceed the currently officially supported 128 virtual machine limit (This has now be increased to 196 offering 170 virtual machines at 1 VCPU per virtual machine)
[2] Compared to figures published by Gartner

Table of Contents

Foreword

TecDem consultants have been providing VMware consultancy in the UK for over 8 years since the arrival of ESX and then VirtualCenter and have designed and installed over 250 installations. The initial two main uses for VMware products were server consolidation and disaster recovery hosting where organisations wished to retain physical servers in production and host replica virtual machines for business continuity.

In 2006 the use of virtual machines to host desktop applications started to emerge and was named Virtual Desktop Infrastructure (VDI). Since the beginning of 2007 this market started to gain momentum as technologies came onto the market that managed the connection of a user's machine (PC or thin client) to a fully functioning Windows XP virtual machine centralised in the datacenter. These solutions, called connection brokers, are the new kids on the block and is the fluid part of the solution stack.

The concept of centralising desktop functions has been around since Citrix and Terminal Server came onto the market over 10 years ago and has been a high priority for most IT departments in all types of organisations. It is a move back to the mainframe days where costs could be controlled if as much of the technology as possible was centralised.

As with most technology solutions the first challenge is to reduce total cost of ownership and the second challenge is to deliver applications and services as well as a desktop can.

How VDI answers the second challenge is by providing a fully functioning Windows XP computer to the user which just happens to be a virtual machine hosted in the datacenter.

How the primary challenge is achieved is by leveraging the cost savings delivered by virtualisation. By taking a large physical server and dividing it up into virtual machines as much as possible is how we deliver TCO savings. This is called the consolidation ratio. The more virtual machines that can be run on a host server the lower the cost per user are achieved.

This white paper specifically discusses consolidation ratios and how high ratios can be achieved using various manufacturer's hardware and types of servers.

All this information detailed in this document is from real world experience from installations we have carried out along with tests performed using hardware we had available.

All numbers, utilisation figures and costs are gleaned from our client's experiences.

Introduction

For eight years we have used an extremely simple calculation to determine the initial savings that can be achieved by installing VMware into a datacenter. It is dividing the cost of the solution by the number of virtual machines to work out a cost per VM. This is then compared to the cost of the machines the client is currently purchasing. When we present at initial meetings we show the following table to demonstrate how the initial hardware acquisition cost savings can be achieved with VMware. This is a calculation for server consolidation ratio.

See the table below

Project to consolidation 25 servers	
Hardware required is two ESX hosts and shared storage, VirtualCenter, network points, rack space, sizing, design and installation services	£40,000
Cost per virtual machine	£1,600
Current cost of new servers	£2,800
Total cost of 25 new servers	£70,000
Total Savings	£30,000
Percentage saving	42.9%

Now if the consolidation ratio can be improved from 25:1 as in the above example to 35:1 the savings are now even greater as shown in the table below.

Project to consolidation 35 servers	
Hardware required is two ESX hosts and shared storage, VirtualCenter, network points, rack space, sizing, design and installation services	£40,000
Cost per virtual machine	£1,142.86
Current cost of new servers	£2,800
Total cost of 35 new servers	£98,000
Total Savings	£58,000
Percentage saving	59.2%

This is a simple calculation and many more factors need to be considered into the TCO calculation. This does show that an increase in the consolidation ratio from 25 to 35 has a dramatic effect on the TCO saving. It should also be pointed out that it is rare any organisation needs to replace all their servers in one go and this needs to be factored into the initial savings. In fact one of the big objections to virtualisation is the initial high price. If someone can buy one server at £2,800 when needed this appears much simpler than a £40,000 investment.

Virtual Desktop Infrastructure Consolidation Ratios

When VDI first started becoming popular it became apparent that consolidation ratios were higher. In the test we ran in January 2006 we took a dual processor, 2.8GHz clock speed with 4GB of RAM as the base ESX host. This machine had a server consolidation ratio of 7:1 when hosting server virtual machines. These virtual machines were migrated to another host and new virtual machines were created with Windows XP, Microsoft Office and Outlook installed on them. We then had users start connecting to the VM's. When this host had server virtual machines on it, the average utilisation of the host was 65% so this became the target for the VDI environment.

We went up to 20 users before the average utilisation of the host machine reached 65%. That was almost a threefold improvement in the consolidation ratio. This machine is an IBM x255 with ESX from VMware and had a total cost of £11,600. This brought the price per virtual machine down to £580.00.

Hardware Types

There are two types of host machines for running VMware. These are Blades and high end physical servers. The terms used are scale out and scale up.

Blade is scale out. They are individually cheaper and there consolidation ratio is lower. Scale up is high end servers which cost more and can host many more virtual machines. We have deployed VMware for server consolidation and VDI implementations on IBM BladeCenter, HP Blade, IBM System X, HP Proliant, Fujitsu Siemens RX800 and Dell 2850 and 6850. The four most popular are IBM Blade and System x and HP blade and Proliant.

We have seen a fairly universal increase in consolidation ratios of 1.6 times when dual core came onto the market with Intel powered machines.

Our goal in every installation is to ensure we help our clients achieve the highest possible consolidation ratios hence achieve the best possible TCO savings.

Cost per user per month model

When looking to deploy a VDI solution we have found most organisations need to see a cost per user per month model that is lower than deploying a standalone desk top for each user. This is similar to the server consolidation cost analysis. Feedback we have received is that if the solution is below the current spend for the same functionality then the sign off process by finance departments is easier to obtain.

With VDI it is usually possible to divert budget from refreshing desktop machines and invest this money into the start of a VDI project.

Consolidation Ratio Calculation

There are four resources that need to be considered when calculating the consolidation ratio. These are CPU utilisation, RAM, network I/O and disk I/O. Disk and network I/O are important when consolidating servers. These resources are not so important with VDI. This is due to the normal drive speed of a desktop machine is 5,400 RPM and network connectivity is usually 100MB. We have performed capacity planning exercises and these have not registered in any project.

CPU and RAM are the primary considerations. For office workers using desktop applications these numbers are fed into a sizing calculator to essentially add up the total processing power and then this is applied to the proposed host machines to determine the specification of the host.

The next step is to ascertain the host machine the client wishes to purchase. Even though we have installed VMware solutions on more HP hardware than any other brand we have not achieved the same consolidation ratios that can be achieved on IBM System x. Specifically the x3850 and x3950. These two models are in a league of their own. Whilst they appear as the most expensive Intel powered servers on the market, by applying the cost per user model they have always delivered the lowest cost per user.

Models from IBM, HP and Dell below these two machines we have seen like for like consolidation ratios and this includes IBM BladeCenter and HP C Class Blade. We have not installed VMware on Dell blades.

IBM V HP Utilisation Test

We were commissioned by a bank in London to test four machines with the specific goal of calculating the consolidation ratio of each model.

The machines to be used for the consolidation ratio tests were:

- Proliant 580 4 socket quad core 2.4GHz processors 64GB RAM
- HP Blade 685 (C Class) 4 socket dual core 2.2 GHz processors 32GB RAM
- HP Blade 680 (C Class) 4 socket quad core 2.4 GHz processors 64GB RAM
- IBM x3850 M2 4 socket quad core 2.93 GHz processors 64GB RAM

The test was to simulate general office type users. A script was recorded of an average user on a PC and installed into a virtual machine.

This was then set up on each host machine and replicated to determine the number of virtual machines each host could run before the target goal of 80% utilisation was reached.

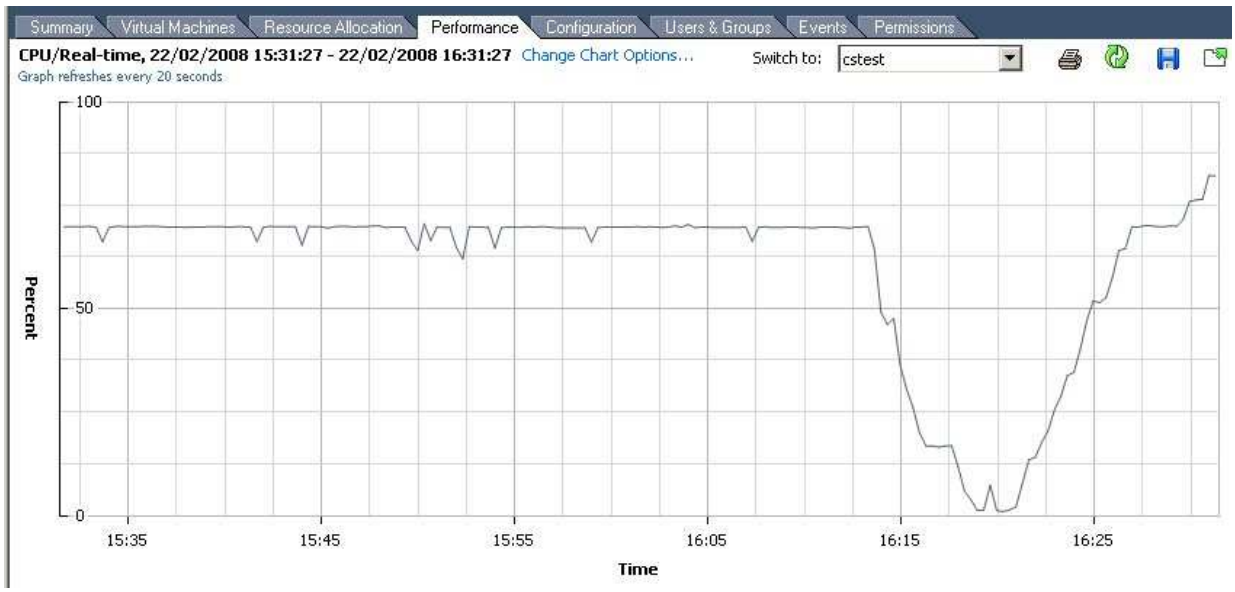

In the above graph the VM's were powered down then back on and the number powered on increased to achieve the 80% sustained host utilisation target.

The utilisation was monitored through VirtualCenter's performance module. Each resource was monitored (CPU, RAM, disk I/O, network I/O) and when any of these resources reached 80% sustained utilisation of the host the test was stopped and the number of virtual machines counted.

The varying processor speeds and amount of memory were taken into account and the final numbers were then extrapolated to give an even number across the board.

It should be noted the HP blades were a single blade in a chassis. This does not provide a real world scenario as a single blade will perform as well as a single server. To this extent we used our sizing calculator to determine the consolidation ratio if the blade enclosure was fully populated.

The operating system installed in the virtual machines was Windows XP SP2.

One quarter of all virtual machines on each host was taken up to a sustained 50% utilisation using an Oracle script to simulate power users similar to a trading environment.

The final consolidation ratios for standard utilisation office virtual machines and ¼ high utilisation virtual machine on each of the host types was:

- Proliant 580 47:1
- HP Blade 685 25:1
- HP Blade 680 27:1
- IBM x3850 M2 98:1

In all cases above the % ready state was also monitored and no virtual machine went above 10%. See below for explanation on % ready state.

Matching CPU and RAM Utilisation

In any virtualisation project the important factor is to match the amount of RAM to the total GHz in the server. As processors get faster and have more cores this has led to a requirement to install more RAM to increase the consolidation ratio.

During a test of virtual machines running Windows XP, MS Office, web based CRM and Outlook we were able to take the utilisation of RAM up to 80% and total CPU utilisation was 41%. This was running 128 virtual machines. Each virtual machine had a script running on it which was recorded over an 8 hour period of a general office worker going about their job.

The specification of the machine was 32GB of RAM and 4 X quad core Intel 2.93Ghz processors. Had this machine had 64GB of RAM installed then we estimate the total consolidation ratio would have been closer to 250. VMware currently limit the number of virtual machines per ESX host at 170 virtual CPU's. Based on this increase from 128 the optimum configuration is 4 X quad core 2.93GHz processors and 48 GB of RAM.

With the increase in allowed virtual machines and on the horizon the increase being higher at possibly 512 a fully specified x3850 M2 would have a consolidation ratio of 250. When the higher increase in allowed virtual machines is made available from VMware the optimum machine would be an 8 X Quad Core x3950 M2 with 128GB of RAM

The above machine with a consolidation ratio of 128:1 would produce a raw cost per virtual machine of £171.88 each. At 170:1 the raw cost per VM reduces to £129.42. This is a 24.7% reduction in costs.

Using VMware's TCO calculator we show below the impact on savings this achieves.

Cumulative 3 Year TCO Comparison	Current (As Is)	With VMware VDI (Projected)	Difference (£ and % savings)
VMware VDI Benefits			
Desktop Capital Expenditure	£ 271,115	£ 0	£ 271,115 ; 100.0%
Desktop Operating Expenses	£ 396,797	£ 186,328	£ 210,469 ; 53.0%
Investment Required			
Capital Investment for Implementing VMware VDI	£ 0	£ 108,705	£ -108,705 ; 0.0%
Total TCO (3 year)	£ 667,912	£ 295,033	£ 372,879 ; 55.8%
TCO average per year per desktop (3 year)	£ 445.27	£ 196.69	£ 248.58

Expected Benefits from VMware VDI	Year 1	Year 2	Year 3	Total
Desktop Capital Expenditure	£ 86,000	£ 90,300	£ 94,815	£ 271,115
Desktop Operating Expenses	£ 64,080	£ 69,976	£ 76,413	£ 210,469
Total Benefits	£ 150,080	£ 160,276	£ 171,228	£ 481,584

Expected Investment in VMware Solution	Initial / Year 1	Year 2	Year 3	Total
Capital Investment for Implementing VMware VDI	£ 36,235	£ 36,235	£ 36,235	£ 108,705

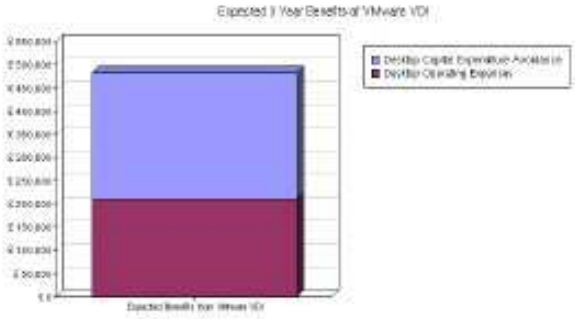

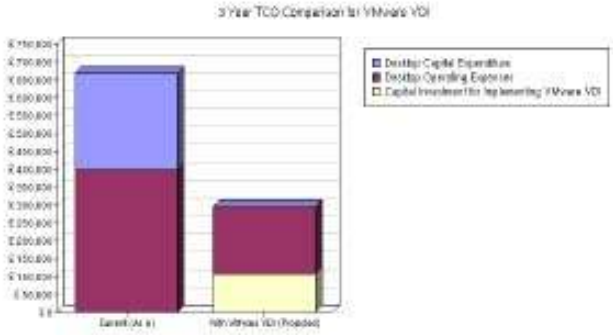

Below is an explanation as to why the IBM x3850 M2 and x3950 M2 achieve such high consolidation ratios.

X4 Architecture

Previous X3 architecture

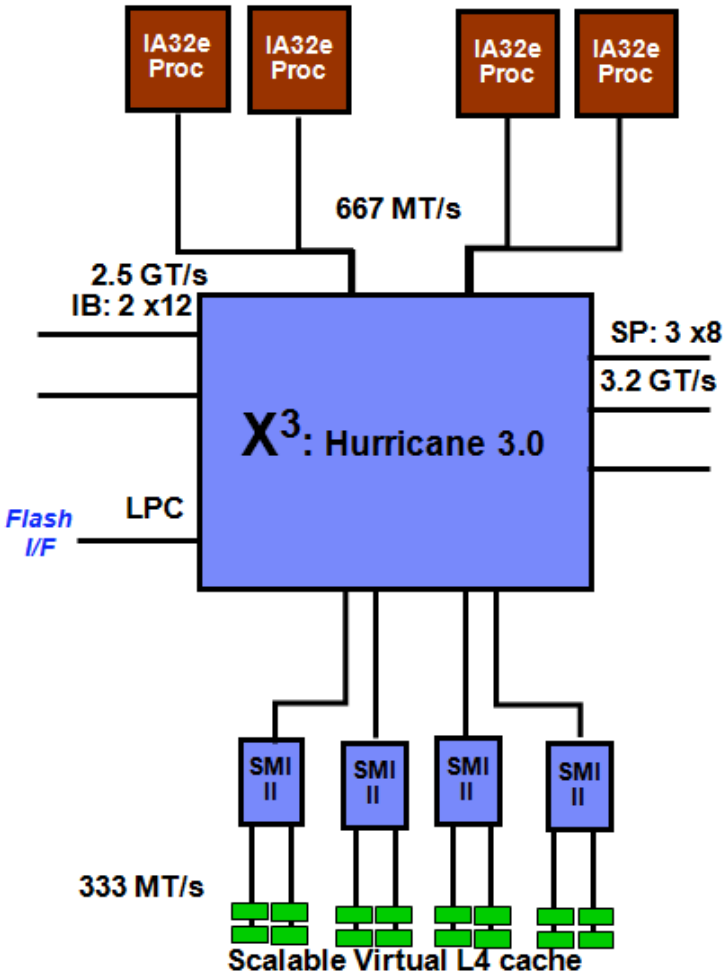

The main part of this diagram shows the number of the channels to the IA32e processors and to the L4 cache. With X3 there was one channel per processor and each pair shared a channel to the controller. With the advent of multi core processors these channels had to provide throughput for either two (dual core) or four (quad core) sets of CPU cycles.

Below is the diagram of X4 architecture. It also shows increases in throughput for the processor channels and the cache.

X4 architecture

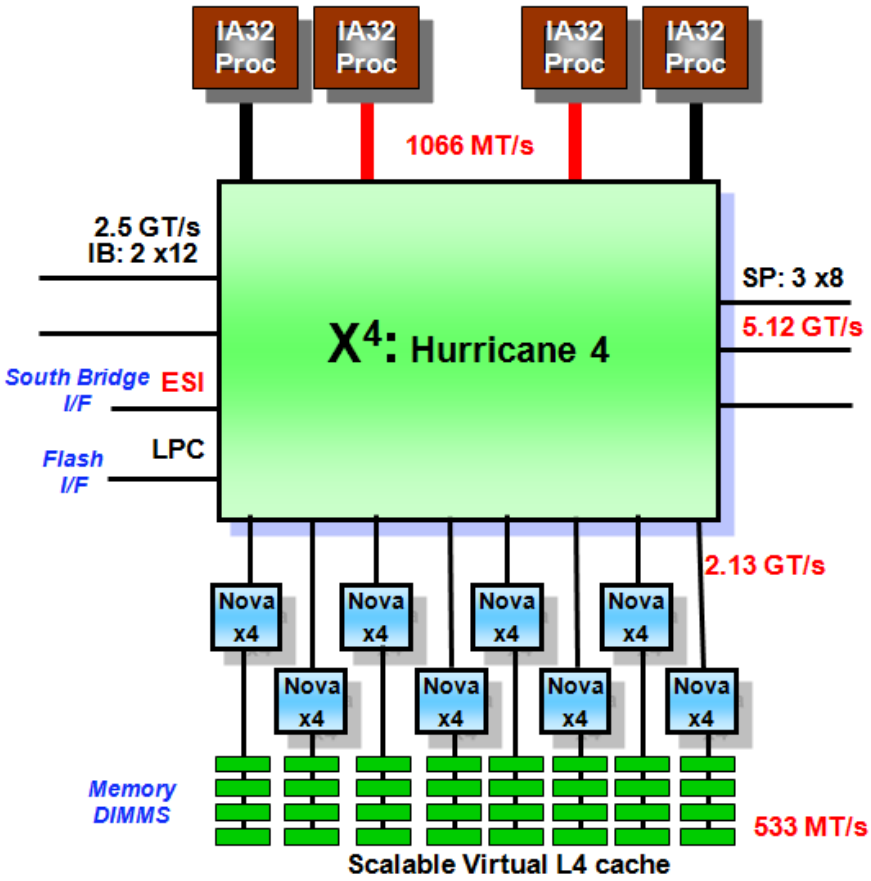

Now each processor has its own channel to the controller, the speed has been increased and each memory set has its own cache.

Our initial tests and user acceptance criteria comparing X3 and X4 machines have shown a doubling of the consolidation ratio using identical virtual machines with load testing software installed in each. When users were asked to connect to their virtual machine they were not informed of any changes.

This is the differentiator between IBM high end servers and any other brand on the market.

% Ready State

The ready state of a virtual CPU is reported from the service console of ESX. This is using the Linux application ESXTOP. When designing a VMware environment and determining the consolidation ratio, the % Ready State should be below 10% for every virtual machine. This tells us that the maximum amount of CPU resources is available for each virtual machine.

Below is an excerpt from VMware ESX Server 3 Ready Time Observations PDF available from www.vmware.com

Interpreting Ready Time

Ready time for a process in isolation cannot be identified as a problem. The best metrics for examining the health of a server continue to be CPU utilization, response time, and application queues.

It is normal for a system to accumulate some ready time even when overall CPU utilization is low. Take an example of two processes (A and B) that each use 20 percent of a CPU, for an overall utilization of 40 percent.

When process B is being scheduled, statistically 80 percent of the time the CPU is idle. The remaining 20 percent of the time process B must wait for process A to finish.

The same is true for process A — 20 percent of the time it must wait for process B to finish. This demonstrates that even under low utilization there is a chance that a shared resource will be busy. Thus some ready time is to be expected and is not a problem. The behaviour is no different in the case of an ESX Server host with multiple running virtual machines. It behaves essentially the same way that an operating system does when trying to run multiple tasks concurrently.

Consolidation Ratio Scenarios

We conducted numerous consolidation ratio exercises to determine the best matches of RAM, CPU, cost of machine and user performance that could be found.

We created two virtual machines. One ran normal office desktop applications including MS Word, MS Excel and MS PowerPoint, web CRM, Outlook and Messenger. The second virtual machine had MS Excel and Oracle trading solution installed in it.

The specification of the virtual machines for the office worker was 1 X VCPU and 512MB of RAM. The trading virtual machine was 2 X VCPU and 1GB of RAM.

The office machine was then set to run at the same utilisation levels as a standard desktop PC which was 12%. The trading system which once started sustained a 50%-55% VCPU utilisation level across both processors. To prove the trading system was truly using 2 VCPU processors at 50%-55% we shut it down and set it to 1 X VCPU. When the application was started it was running at 100% utilisation. It was then switched back for the purposes of the consolidation ratio tests.

Office Worker Consolidation Ratio

The virtual machine was converted into a template and using the auto deployment feature in the connection broker we used for this test (Virtual Access Suite) we started creating virtual machines up to 128. (We did do a 129th machine just to see if VMware would let us go further but it would not power on saying the maximum amount of virtual machines had been reached.)

We then started the scripts on the 128 virtual machines to simulate using the applications installed on the virtual machines. The total utilisation of the IBM X3850 M2 never went above 41% CPU utilisation (across 4 X quad core 2.93GHZ processors) and 81% RAM utilisation. Disk I/O and Network I/O were minimal as we have seen with all VDI implementations on high end servers to date.

Based on this scenario the ideal machine for this type of worker is a 2 X quad core machine with 32GB of RAM whilst VMware keeps the limit to 128 virtual machines per host.

We were not able to try the same test with XEN Source which we are informed does not have an upper limit on the number of virtual machines. We intend to run the exact same tests in the future using XEN Source.

Trading Solution Consolidation Ratio

The same process was performed with the trading solution virtual machine. Knowing it would not get to 128 we created virtual machines in batches of 10. The application was started on each virtual machine and left to push up utilisation.

The theory on this test is we would get exactly 32 virtual machines before all CPU time was consumed. This is based on each virtual machine taking 50% of 2 cores. The host machine had 16 cores so 32 would be the limit before degradation occurred for the trading application.

At 32 virtual machines the machine had reached 64% CPU utilisation. RAM was also matching CPU at 62% showing that for high end applications the 4 X quad core configuration is more suitable.

We then deployed another 14 virtual machines and monitored the % Ready State to see when degradation of performance would kick in. At 46 virtual machines the % Ready State rose above 10% to 12% for the top 3 virtual machines. The host was running at 90% utilisation across CPU and RAM. See graph below.

Whilst we had this test running we decided to test a marketing message from VMware. They state that ESX will allow more virtual machines to run on a host than appears physically possible. This is called over subscribing the machine and useful for reducing DR costs.

We added another 10 virtual machines and hit the 100% mark. Each virtual machine kept running proving it is possible to over subscribe. However we don't recommend this for production.

See the graph below

During the 128:1 consolidation ratio exercise the highest % Ready State was 9% running down to 4%. This is considered well within tolerance levels.

During the 46:1 consolidation ratio exercise the highest % Ready State was 12% running down to 7%.

Power Consumption Test

As many organisations need to reduce power consumption due to lack of available power into datacenters and also to reduce CO_2 emissions we ran the same tests on an IBM x3850 M2 to measure power consumption.

For this test VMware allowed us to go above the 128 active virtual machine limit.

The machine had a power clamp connected to a power board and both power supplies were connected to the board to ensure all consumption was measured.

The test was designed to calculate the amount of wattage per virtual machine. A second test was also carried out by taking the utilisation of virtual machines up to 100% utilisation using CPUBURN.

The computer room temperature was maintained and monitored at 19°C during the tests.

First Test

The host was isolated from the other machines in the cluster. VirtualCenter, Active Directory, 2 X file servers plus the test virtual machines were retained on the host to perform the test.

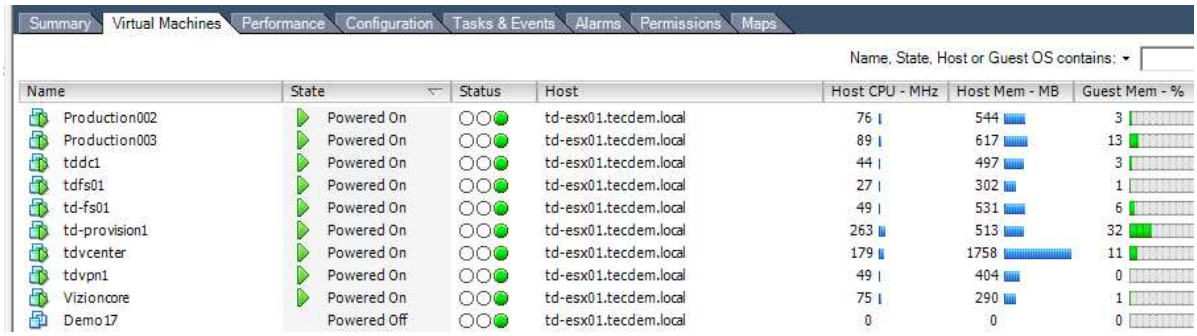

Name	State	Status	Host	Host CPU - MHz	Host Mem - MB	Guest Mem - %
Production002	Powered On		td-esx01.tecdem.local	76	544	3
Production003	Powered On		td-esx01.tecdem.local	89	617	13
tddc1	Powered On		td-esx01.tecdem.local	44	497	3
tdfs01	Powered On		td-esx01.tecdem.local	27	302	1
td-fs01	Powered On		td-esx01.tecdem.local	49	531	6
td-provision1	Powered On		td-esx01.tecdem.local	263	513	32
tdvcenter	Powered On		td-esx01.tecdem.local	179	1758	11
tdvpn1	Powered On		td-esx01.tecdem.local	49	404	0
Vizioncore	Powered On		td-esx01.tecdem.local	75	290	1
Demo17	Powered Off		td-esx01.tecdem.local	0	0	0

With these virtual machines running the utilisation was an average of 5.34%. This was the base of an ESX host at idle.

Virtual machines were then powered on in increments of 20 and the utilisation scripts started to simulate office workers. Power consumption of the host was measured after 15 minutes on each increment to allow the host to settle from higher utilisation that is caused by starting virtual machines.

Below is the table showing the amount of overall power consumption of the host during the incremental increase in running virtual machines and host utilisation.

Qty of VM's	Overall watt usage	Watts per VM	Overall % consumption increase	% reduction of consumption per VM
0	657	N/A	N/A	N/A
20	682	34.10	1.04%	N/A
40	713	17.83	1.05%	-47.73%
50	745	14.90	1.04%	-56.30%
60	798	13.30	1.07%	-61.00%
80	822	10.28	1.03%	-69.87%
100	838	8.38	1.02%	-75.43%
120	841	7.01	1.00%	-79.45%
140	842	6.01	1.00%	-82.36%
160	844	5.28	1.00%	-84.53%
180	849	4.72	1.01%	-86.17%
200	861	4.31	1.01%	-87.38%

Overall % consumption increase shows a linear increase of just over 1% for each increment with a total increase in power consumption from idle to 200 of 23.7%.

At 120 virtual machines the watt usage per VM of 7.01 represents an excellent power cost per user.

We then tested a thin client and monitor to calculate an overall watt per user number. The thin client was a Wyse V90L and the monitor was a Phillips 17" LCD. The total consumption was 44 watts. Combined with a consolidation ratio of 120:1 and an allowance of 5 watts per user for storage the total power consumption equates to 56.01 watts.

It is not possible to calculate a watt per user on storage as this is so variable. For the purposes of this calculation we used an IBM DS4300 with 1.4TB of storage. There was less than 20 watts variance from when the virtual machines were powered off and on. Its total power consumption was 726 watts. The estimate of 5 watts was used. The SAN could easily accommodate 4 X the amount of users as a host so this estimate is on the high side.

With the increase of virtual machines allowed by VMware the watts used per user drops by .49 which does not show a marked improvement in power savings as this is already achieved at the 120 mark.

Test Two

With all this equipment in place we decided to take the host to maximum CPU and RAM usage to see how much power the host would draw. The x3850 M2 has dual 1,300 watt power supplies.

We loaded CPUBURN onto 46 virtual machines and kicked this application off. The overall utilisation of the host rose to 91%.

Name	State	Status	% CPU	% Memory	Memory Size - MB
td-esx02.tecdem.l...	Not responding		1	4	32767.55
td-esx01.tecdem.l...	Connected		91	80	32767.55

Average watt consumption during this test was 923 watts. (We don't recommend taking a production machine up to this limit)

Further savings in power consumption can be achieved with the new feature in version 3.5 of ESX and 2.5 of VirtualCenter which is called Distributed Power Management. (DPM).

DPM is still experimental from VMware. This feature monitors the overall utilisation of a cluster of ESX hosts and when utilisation drops below a set threshold DPM will automatically VMotion virtual machines away from hosts so they can be powered down. This is integrated with DRS so load balancing takes place on the remaining hosts. In theory a cluster of 10 hosts running VDI virtual machines could power down 8 hosts at night as users log off and go home. Hosts would then be powered on in the morning to ensure there was sufficient resources.

Proof of Concept Options

As with just about all VMware/virtualisation projects organisations wish to run a proof of concept to determine TCO savings, ROI, user acceptance, application compatibility and suitability of the solution for their own environment.

The ideal process for building a successful POC is to first ascertain the specification of the desktop machines, types of applications and average utilisation. Another key factor is the percentage of desktops that are actually used at any one time. On a TCO calculation project we carried out last year for a company with 1,200 users we found that the highest amount of active PC's at any time was 620. Had this client designed a solution to host 1,200 VM's it would have run below 40% utilisation. In the end the final design was for enough hosting capacity for 750 users to allow for headroom. By actively monitoring the utilisation of ESX hosts it is easy to add more capacity as required for user count increases. The initial cost for 1,200 users was just over £1M. Reducing the capacity to 750 cut £250K of the bill providing a faster ROI.

There are tools on the market to analyse utilisation levels of PC's however in our experience this is an unnecessary cost. Usually the simplest method is to add up the target machines, adding the total CPU GHz and RAM then determining the type and specification of the server that would be the ideal host.

Once this information is complied a POC can be put in place and using the same tests carried out to create this white paper it is possible to calculate the consolidation ratio. Once the POC machine is installed the utilisation numbers can be extrapolated out to many users providing an initial guide to the cost of the implementation. Divide this by the number of virtual machines and you have a cost per VM price.

The ideal POC should be isolated from any existing server virtualisation implementations so accurate utilisation numbers can be obtained.