How fast is the cloud?

intechnica

Digital Performance

Contents

Introduction

This research project has been undertaken by Intechnica for two clear reasons: firstly the cloud matters. Gartner predicts that by 2015 50 per cent of CIOs expect to operate their applications and infrastructure using the cloud. IT leaders need to be equipped with reliable and robust data in order to make informed purchasing decisions.

Secondly, performance matters. End users are demanding faster and faster response times. 57 per cent of online consumers will abandon a site after waiting three seconds for a site to load and 80 per cent will not return. For companies with revenue-generating websites, there is a direct correlation between performance and revenue. If a site is not performing, then it is losing money. More and more business critical functions are being carried out online; the impact of poor performance or failure of these systems could be catastrophic. In addition, web applications are becoming increasingly complex and rich, so delivering performance is getting simultaneously more important and more difficult to achieve.

What is the cloud?

There are many statements, claims, counter-claims and promises made about the cloud. Getting to the truth is not easy, and seems to be getting harder as the increasing hype, media attention, and 'expert' opinion grows day by day. We believe the quotation below is a nice place from which to build a clear understanding of the cloud.

"Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interaction."

The National Institute of Standards and Technology

"If, as analysts predict, we will see an annual double-digit growth for the next 5 years at least a clear understanding of the cloud is vital for businesses and specifically for IT leaders.

For the purposes of this study we are looking at an ecommerce application installed on Platform as a Service (PaaS), and Infrastructure as a Service (IaaS) offerings.

We are not examining any Software as a Service (SaaS) offerings though a similar system could be created using platforms such as eBay or Amazon Marketplace".

2015

The year when 50% of all CIOs expect to operate the majority of their applications and infrastructures by the cloud

(Gartner)

http://www.gartner.com/it/page.jsp?id=871113

An enabling technology

For decades the IT leader has been split between how they can add strategic value to a business through technology and the tactical day to day running of a reliable IT system.

It was often a thankless task for CIOs who found themselves struggling to juggle the two. However, with cloud computing CIOs can 'rent' the technology, potentially reducing costs and management overheads, whilst gaining greater flexibility and autonomy. It can enable greater risk taking, allowing new approaches, environments and applications to be created, trialled and taken down again in a matter of days, and at much less cost than on conventional infrastructures.

This frees up the CIO to add real value to a business and concentrate on strategic management issues that will save money or – even better – drive revenue growth.

As well as taking away some of the challenges of running an effective IT resource, the cloud can be used to achieve competitive advantage, to create efficiency, stimulate collaboration and facilitate easier information sharing.

The challenge is how best to use the cloud. Our research aims to focus on just one area: that of performance. To keep things in the real world we have also calculated total cost of ownership for each platform. We have sought to provide a clear, expert assessment of the performance of cloud in order to help IT leaders make better informed purchasing decisions, and to contribute to the body of public knowledge and information about the cloud in general.




Executive summary

The single most important outcome of our research is that the cloud can be at least as fast as conventional hosting environments, and when the total cost of ownership is considered, it becomes an extremely compelling solution that businesses can not ignore. In fact, our project should reassure any IT leader who has been cautious about adopting the cloud because of uncertainty about performance. However, what our research proves is that good performance comes by design, planning and good management, not by accident.

In our tests, Microsoft Azure and our VMWare public cloud both performed well, whilst Amazon EC2 (AWS) performed the least well of our four environments. This in no way suggests that AWS is the least 'good' cloud option; in different test conditions using a different application, AWS is quite likely to outperform the rest of the group. In fact, with some minor changes to the environment taking just a few hours, we were able to improve AWS results by 57%.

In addition to performance testing the cloud, we also calculated total cost of ownership for each of our 4 environments. The figures reveal some interesting insights into how the cloud can be considerably more expensive than physical environments if not used intelligently.

In summary, our results provide an excellent illustration of several fundamental truths about moving to the cloud that any business should be aware of:

1. The cloud is different to conventional physical hosting environments, and it requires intelligent upfront and ongoing management to maximise the benefits and avoid the pitfalls

2. Not all cloud solutions are the same; consequently, they need to be considered separately and dealt with differently to achieve optimum performance.

3. Applications need to be designed or modified specifically for the cloud – simply migrating conventionally hosted applications is likely to result in failure.

4. The cloud is a liberating and enabling technology: businesses can have confidence to conduct proofs of concept, and take greater risks with technology choices without the fear of large costs, significant deployment of resources and long-term commitments to third parties.

Top 10

Do's and Don'ts to achieve cost-effective performance in the cloud:

- Do adopt the cloud – the benefits far outweigh the downsides.
- Don't treat the cloud the same as conventional hosting environments.
- Don't think of all cloud solutions as the same.
- Don't migrate applications to the cloud without properly assessing their suitability beforehand. Do research and plan ahead – know the pros and cons, the benefits and pitfalls before you launch into it.
- Do design and build applications to be suitable for the cloud.
- Do consider hybrid options, for example where sensitive data remains in a secure environment, but front-end processes are carried out in the cloud.
- Do a proof of concept for the cloud, choosing low risk, short term, discrete projects that are likely to benefit from the main benefits of flexibility, scalability.
- Do establish a different set of internal working practices and disciplines to ensure you maximise the benefits and avoid the pitfalls on an ongoing basis.
- Do put in place tools and processes to monitor performance
- Don't assume your knowledge is transferable; consider engaging independent consultants before selecting your cloud platform and before migrating your applications.

Project overview

In this report we provide an overview of our project approach and findings. A more detailed description and analysis is available in the Appendix.

With this research project, we set out to simulate the experience of a buyer seeking to exploit the benefits of the cloud, and specifically to challenge the assertion that the cloud delivers worse performance than conventional hosting options.

We wanted to test, within a real-world situation, the relative performance of two of the largest and best known public cloud offerings (Microsoft Azure and Amazon AWS), a VMWare public cloud (in this case supplied by VirtualDCS), and a conventional physical server environment.

It should be noted that we have limited the scope of our research purely to using the cloud for the hosting of public-facing websites, and not to the many other uses that the cloud can be considered for (including, batch processing, development environments, internal file servers, etc).

Our approach

We installed an open source, ecommerce application onto four different infrastructures: Azure, AWS, VMWare and physical servers.

The application was chosen, since it represented the two-tier application infrastructure common in many business applications. Also, as the application was built on the .NET framework, and had already been ported to Azure, it meant that it was supported by all of our target infrastructures.

The specification of each environment was kept as similar as possible, and we ensured they remained largely "out of the box".

We ran a series of performance tests, using Intechnica's **TrafficSpike** tool to simulate key business transactions, including a mix of page browsing, searching and transactional processes.

We used a series of standard monitoring techniques, analysing page response times as well as key system counters relating to the performance of the web and database servers. See the Appendix for further detail on the research methods and findings.

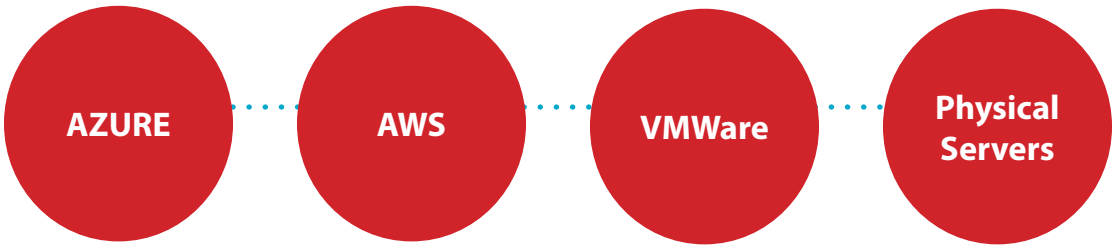

What did we find out?

Our tests supplied some fascinating and surprising results for the cloud. The results for Azure and VMWare demonstrate unequivocally that the cloud can perform as well as, if not better than physical servers.

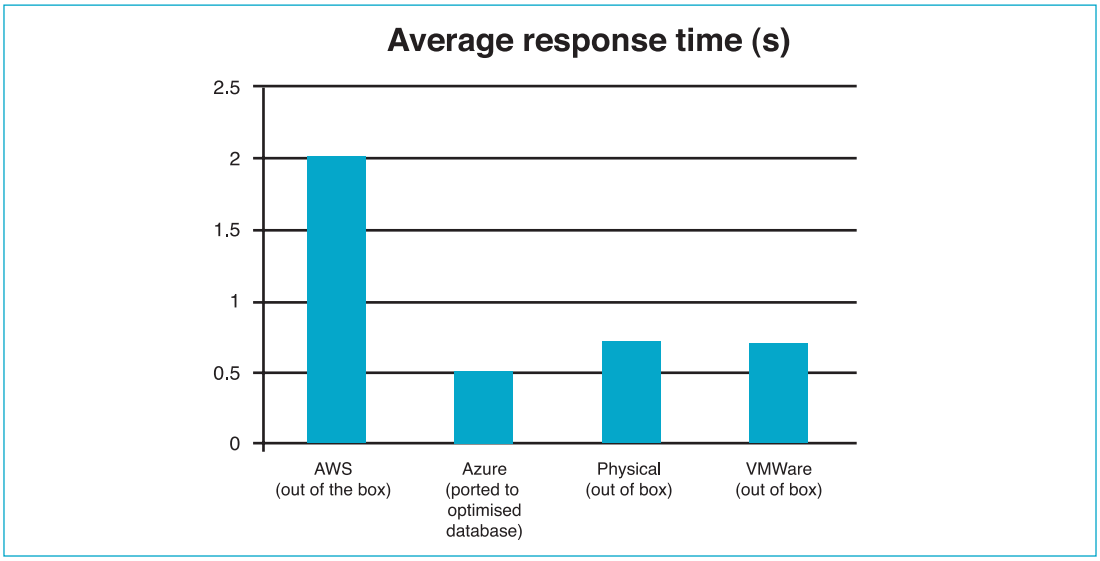

If our research serves no other purpose than to have dispelled the myth that the cloud is not fast then it has been a worthwhile exercise. However, it does this and much more as the results for AWS show: sometimes, in certain circumstances, some cloud solutions will perform significantly worse than physical servers and indeed other clouds.

This finding is fascinating and bears out a couple of fundamental truths:

1. There is no performance hit directly related to virtualisation. Physical kit and VMWare had very similar specifications and provided near identical results.

2. Public cloud services (such as Azure and AWS) are not equivalent to VMWare solutions

3. Applications optimised for cloud platforms and taking advantage of highly scalalable cloud services will outperform physical kit.

4. Not all clouds are the same. Your choice of cloud option should be made after careful and expert consideration of many factors including: your commercial and IT department cultures, priorities, budgets and attitude to risk; your existing applications and IT estate; and the business processes and end-user behaviour that will be dependent on your choice of cloud.

So, what happened to AWS? Why did it perform the least well of the four environments?

The application that we selected for our research is very database-dependent. As well as customer and product data, even product images are stored within the database.

Quite simply, AWS and the application we used for the test are not a good match. After identifying the database as a potential bottleneck, we made some minor changes to commonly-used SQL stored procedures. These changes, which took a matter of hours resulted in a 57% performance improvement.

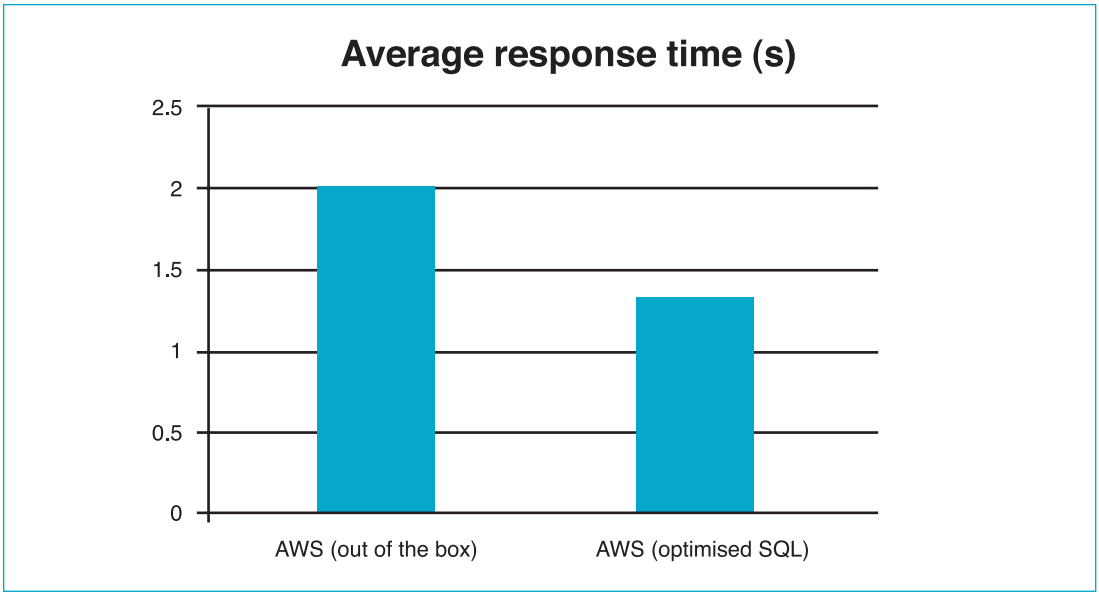

Whilst this in itself has shown an impressive leap in the results, this is just the start of the story, and the performance can be driven further still.

Storing images within the Amazon Content Delivery Network (CDN) would improve performance in two ways. Firstly by exploiting improved caching and download speeds from CDN and secondly by reducing the load on the database server.

An alternative back-end database such as Amazon Relational Database Service (RDS), which is optimised for this platform, could be used to further improve application performance.

Total cost of ownership

To fit in with our aim of simulating a real-world situation, we decided it was important to compare the relative costs of each option. If the best performing option was significantly more expensive, or indeed significantly cheaper than other options this would obviously impact purchasing decisions.

Comparing prices is a complicated process as each platform uses its own charging structure, weighting each aspect of their services differently.

"To establish TCO, costs were assessed based on the likely annual price of set up and operation using the app on each platform. Items included the cost of the server, set-up, storage, SQL, database usage, bandwidth and administration"

Monthly costs for the two servers used in our tests

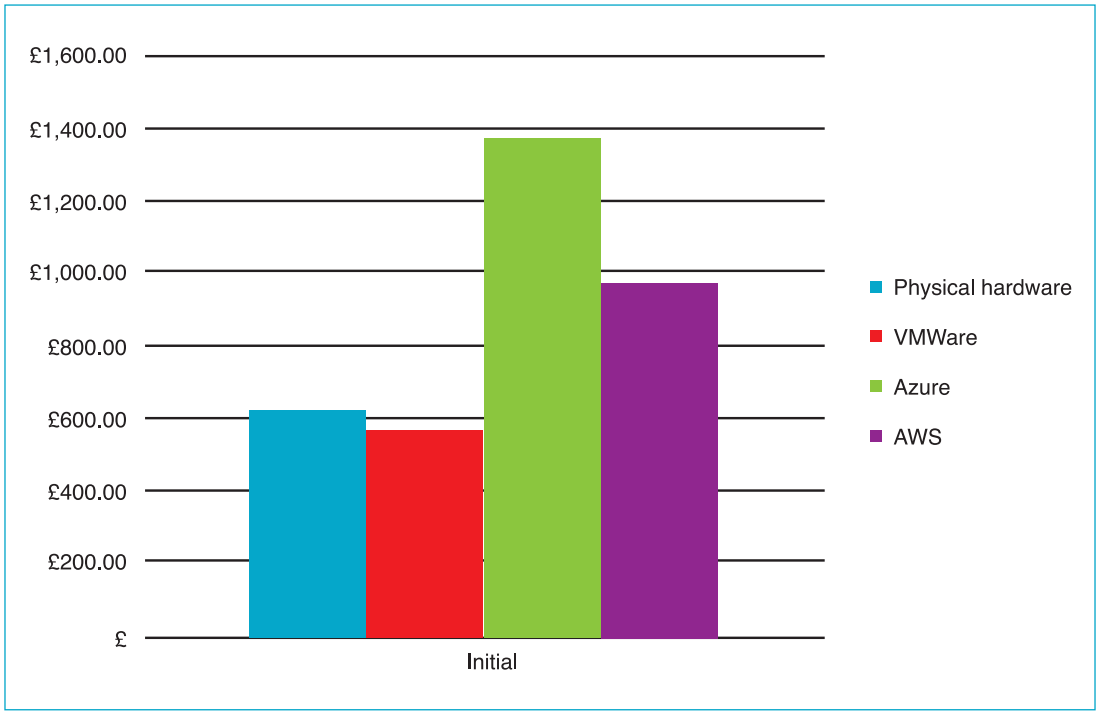

The above chart demonstrates that for our "artificial" pricing model, based on two servers in use 24 hours per day there are stark differences in the costs which are not flattering to AWS and Azure.
(See further breakdown of costs in Appendix 2).

Only when using cloud options sensibly can cost-benefits be realised, for example in a real-world situation test environments, development environments and Disaster Recovery (DR) environments all contribute towards IT costs. Many of these environments are used infrequently, if at all.

If we re-calculate our costs based on development environments used

30%

of the time and performance test environments used

10%

of the time, significant cost savings can be demonstrated.

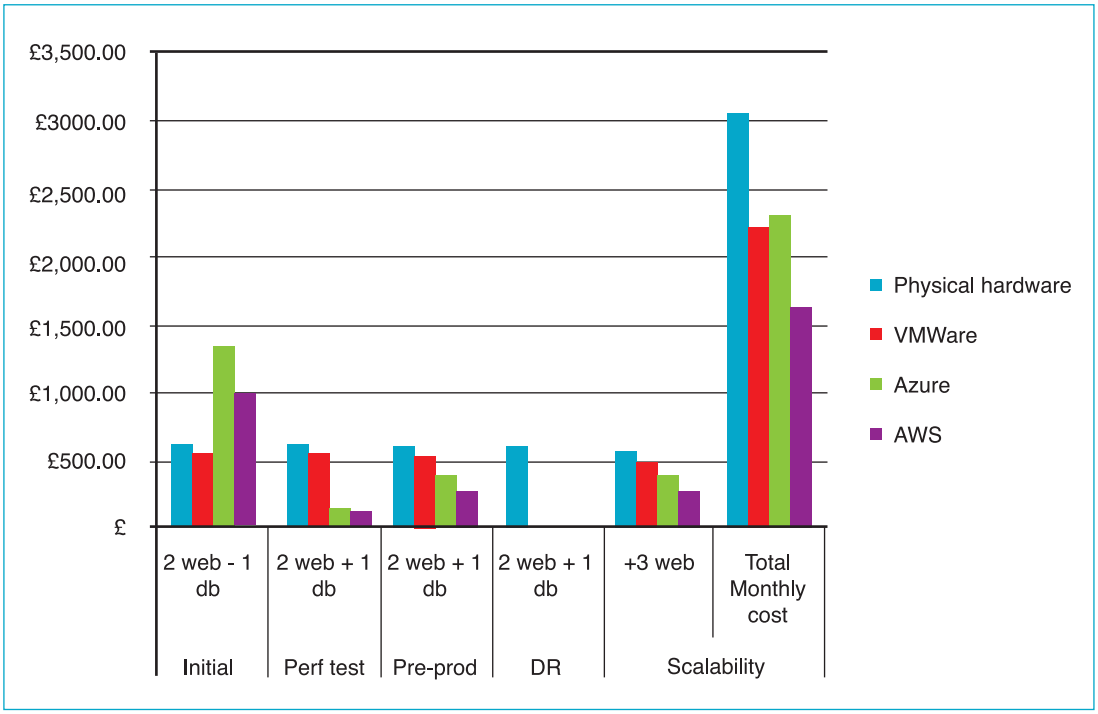

Only paying for what you use had significant cost benefits when comparing VMWare, Azure and AWS with traditional hosting.

In summary, as the diagram below demonstrates, the cloud demands more sophisticated and intelligent use in order to realise the significant cost savings available. Moreover, an ill-thought through or poorly managed use of the cloud will result in painful financial consequences.

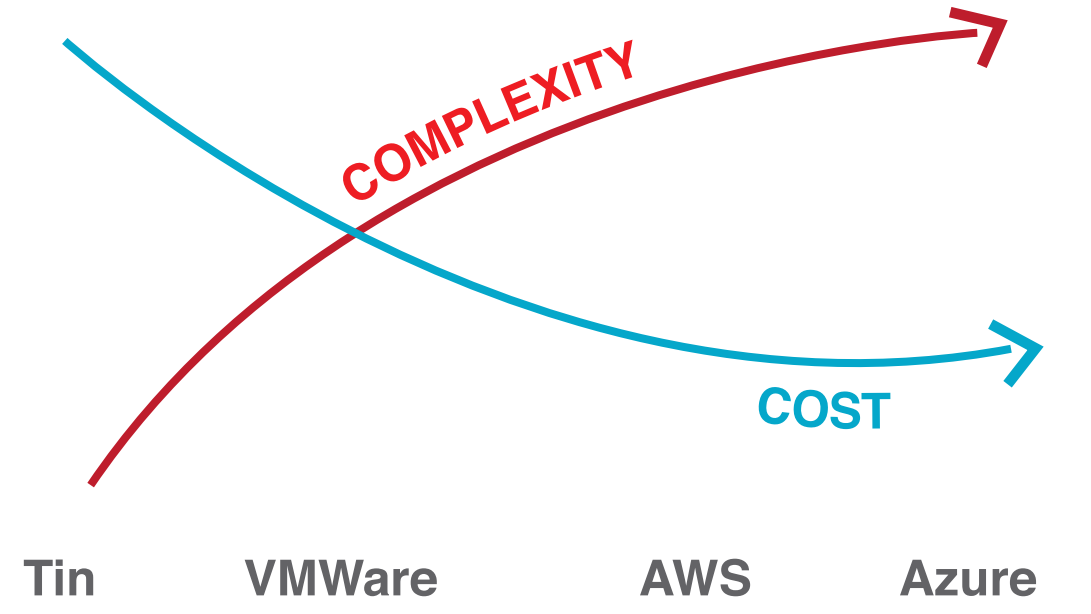

Complexity in this chart refers to intelligent management and usage: processes and disciplines for effective use of the cloud need to be established and managed (e.g. turning off instances when not in use).

It does NOT refer to technical complexity.

Analysis

What does all this mean?

1. Cloud based applications **can** compete in terms of performance but **won't always** compete

The results provided by Azure show that a standard eCommerce application can out perform the same application using physical hardware. The results shown by AWS illustrate that this is not a given.

This is not a surprising result when you consider some of the services that run on this architecture.

2. Cloud development ≠ traditional web development

Don't believe that if you can develop for the web you can develop for the cloud. Don't believe a web application will transfer seamlessly to the cloud. It may do but you should be prepared to experiment.

Developing for the cloud involves a mind shift to take advantage of the benefits offered by the cloud but requires a constant awareness of the weaknesses. These can mean performance issues but can also mean subtle failures of the application.

3. Flexibility can be a weakness in terms of performance

We see this with AWS in this test – AWS is the most flexible of all the platforms we tested.

Because of the wide variety of platforms available and the ability to install any software needed it is easy to get an application up and running. However this will not necessarily perform the same as a physical environment.

4. Be prepared to experiment with cloud implementations

Cloud services are ideal for experimenting with what does and what doesn't work. Build your system on one platform, test it and if it doesn't work tweak with the installations, machine size, platform options. All this can be done with no hardware investment. If it doesn't work try another cloud platform or adapt your application until it does.

Developers traditionally do this with technologies and languages in the development environment, with cloud platforms the same luxury can be afforded to the deployment environment.

5. To get the best performance from the cloud make the most of the services offered by cloud providers

The big cloud providers have created services aimed at high usage. These include highly scalable SQL and NoSQL databases, shared storage, access layers, caching and more. Using these will make your cloud based application perform better.

This is borne out in the test results where we see Azure, using SQL Azure out performing AWS which is using an EC2 based install of SQL Server rather than Amazon's RDS service. Amazon's first recommendation for improving the performance result – use RDS rather than SQL Server.

Cloud providers are experts on scalability and the services they provide are based on those used to power their own system. So Amazon SimpleDB is the same technology used to power the Amazon website.

Cloud providers are experts on scalability and the services they provide are based on those used to power their own system.

So Amazon SimpleDB is the same technology used to power the Amazon website.

At a glance

The four platforms

PHYSICAL SERVER	
PROs	**CONs**
Mature service industry for support of platform	No flexibility, once bought, need to keep.
Highly configurable for scale-up solutions	Limited ability to scale up or scale down rapidly to accommodate seasonal peaks in demand. Can buy more kit or upgrade, but not quickly.
Good performance – bang for buck	Ongoing support and maintenance costs
Financial certainty, write down over e.g. three years	Responsible for operational management
Multiple OS available	

VMWare hosting co. (IaaS)	
PROs	**CONs**
Mature service industry for support of platform	Short term contracts not common
Good performance (comparable with physical hardware)	Similar support and maintenance overhead to physical server
Some benefits from virtualisation, shared resources etc.	May be responsible for operational management
Shorter term options than outright purchase	
Multiple OS available	
Limited ability to scale out and scale up (more so than physical server, less than AWS/Azure)	

AWS (IaaS)	
PROs	**CONs**
Similar support methods to Physical / VM Multiple OS available Pay for what you want when you want it Scale out easily Very short term options available for complex environments making it easy to build short term test environments or implement short-term business solutions Short lead time for provisioning instances Access to CDN functionality Enterprise core components, e.g network infrastructure / backup / resilience	Less certainty over price model Comparatively less bang for buck compared with physical server/VMWare Not as simple to manage as you might think (similar to server)

Azure (PaaS)	
PROs	**CONs**
All patching / updates done for you – low management overhead Simple upgrade path from small to large instances. Common 64-bit platform from "extra small" to "extra large" instances, simplifies route to live for developers. Very short term options available for complex environments Pay for what you want when you want it Scale out easily Very short term options available for complex environments making it easy to build short term test environments or implement short-term business solutions Short lead time for provisioning instances Access to CDN functionality Enterprise core components, e.g network infrastructure / backup / resilience	Only one choice of OS Less certainty over price model Comparatively less bang for buck compared with physical server / VMWare "Learning curve" compared to IaaS Apps need to be written for Azure or ported over Support model immature

About Intechnica

Digital Performance

Intechnica are a digital consultancy specialising in online application development, performance services and cloud consultancy. Our clients include:

Performance Services

Performance services from Intechnica can help discover what could go wrong with your business critical system before it actually does. We start by understanding our client's business which informs everything we do and determines the range of services, tools and resources that we deploy.

TrafficSpike

TrafficSpike is Intechnica's on demand load testing service offering multiple different traffic injection methods, including from the Amazon EC2 cloud environment.

dynaTrace software

dynaTrace are a technology leader in Application Performance Monitoring offering a solution which provides a far greater depth of monitoring than that previously available in a monitoring tool.

Application Development

Our unique development methodology "Automated Test Driven Development" builds performance, quality, scalability and agility into your application from day one. If you have ever been told by other developers that you can't have quality, time and low costs then our approach could be for you.

We believe in combining the best people with the best tools to produce high quality, innovative solutions which are built to last and designed for growth.

Cloud Consultancy

Intechnica can help you harness the power of the latest cloud technologies to reduce capital expenditure and utilise the benefits of flexible computing resourcing. To maximise the benefits of the cloud it is vital to design applications to use it:

- Design for scale-out to make appropriate use of cloud resources
- Design to allow separate business processes to operate in isolation
- Design with cloud pricing models in mind to minimise costs

Appendix

Application under test (Target application)

nopCommerce - an open source e-commerce application based on ASP.NET 4.0 based with a MS SQL backend database.

Reasons for choosing application

We chose nopCommerce because it was a good example of a two-tier .NET based application in widespread use. The application had already been "ported" to the Azure platform and a simple installation could be carried out onto each of the other target platforms.

Demo site:

http://demo.nopcommerce.com/default.aspx

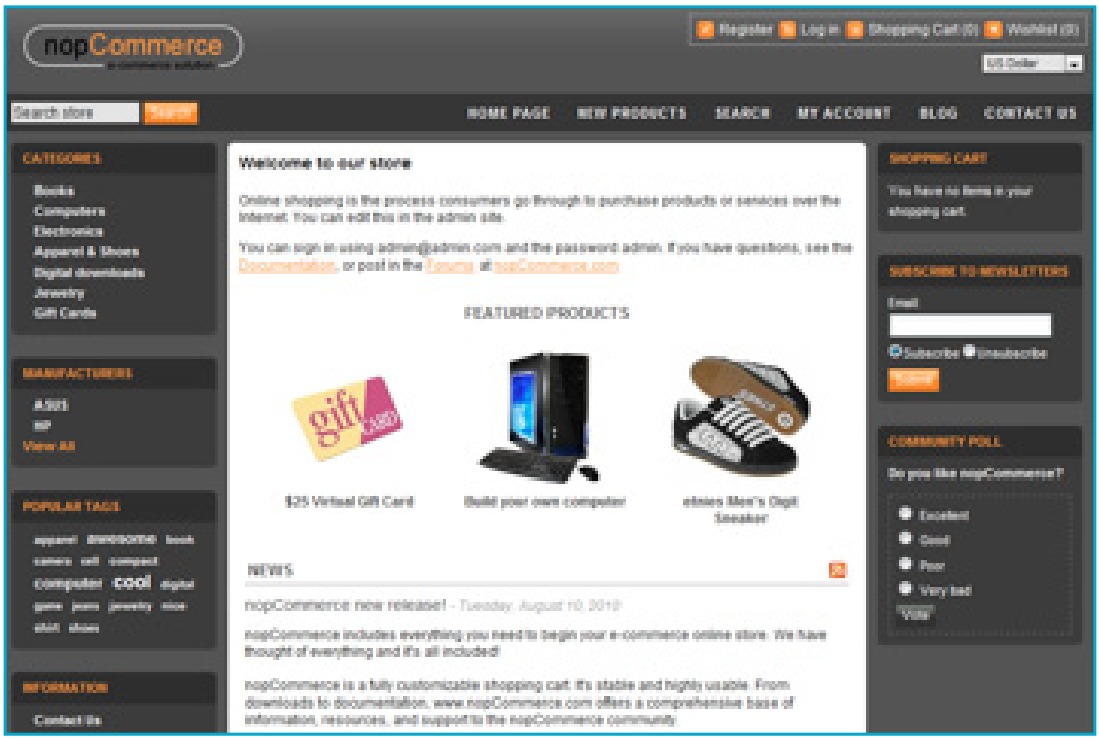

Test Approach

We developed scripts using our cloud-based performance test tool, **TrafficSpike**. This allows us to simulate a load on the test platforms via the Internet to compare page response times for business transactions.

TrafficSpike

Test scripts

We developed test scripts which simulate the following business processes: Browsing for products, Searching for products, Adding items to shopping cart / wish list, Editing Wish list / cart, checking out and customer registration.

Test scenarios

To simulate business-like load we set proportions of virtual users as follows:
85% of users performing browsing / searching activity
15% of users purchasing items and registering with the site.

Tests were configured to "ramp up" to peak load over a 5-minute period, maintain load for 30 minutes and then ramp-down.

Load equivalent to 46 users pausing for between 5-10 seconds between page requests was simulated. Approximately 10,000 page requests were made during the 30-minute period at peak load.

85% of users performing browsing / searching activity

15% of users purchasing items and registering with the site

Tests were configured to "ramp up" to peak load over a 5-minute period, maintain load for 30 minutes and then ramp-down

Load equivalent to 46 users pausing for between 5-10 seconds was simulated

Approximately 10,000 page requests were made during the 30-minute period at peak load

Load profile:

Transaction Analysis screenshot:

Start Page | Runs Explorer: LoadTest | Run Summary | VU Concurrency | Group Summary | View Summary | View Summary

LoadTest Baseload1.27::Plateau::View Summary

View Summary | Transactions | Transaction Percentiles | Web Requests

Transaction	Count	Maximum	Mean	Minimum	Errors	Median	95th Percentile
<COMBINED>	9910	5.446	0.712	0.060	0	0.489	2.355
AddToCart	54	3.172	1.156	0.511	0	1.073	1.741
AddToWishlist	89	3.437	0.898	0.188	0	0.792	1.602
BrowseToProducts_Accessories	41	2.820	0.816	0.449	0	0.582	1.695
BrowseToProducts_Apparel & Shoes	85	3.283	0.811	0.446	0	0.599	1.644
BrowseToProducts_Books	86	2.099	0.678	0.408	0	0.528	1.341
BrowseToProducts_Cell phones	85	1.637	0.719	0.435	0	0.650	1.192
BrowseToProducts_Computers	84	2.804	0.847	0.466	0	0.621	1.741
BrowseToProducts_Desktops	13	0.921	0.607	0.428	0	0.597	0.921
BrowseToProducts_Digital downloads	75	1.469	0.557	0.372	0	0.460	0.933
BrowseToProducts_Electronics	85	2.163	0.724	0.445	0	0.519	1.715
BrowseToProducts_Games	20	1.827	0.805	0.460	0	0.540	1.809
BrowseToProducts_Gift Cards	86	1.771	0.750	0.376	0	0.603	1.365
BrowseToProducts_Jeans	22	1.463	0.703	0.460	0	0.587	1.223
BrowseToProducts_Jewelry	80	1.621	0.615	0.370	0	0.480	1.359
BrowseToProducts_Notebooks	18	0.984	0.570	0.423	0	0.484	0.888
BrowseToProducts_Shirts	21	0.996	0.533	0.404	0	0.464	0.837

Test data

the nopCommerce application has less than 100 products in the database by default for testing purposes. We increased this to 7,500 products each of which had multiple images associated with it. Text for product names and product descriptions was extracted from a copy of "War and Peace" downloaded from the Project Gutenberg archive.

Products were created in bulk using an SQL script which created products, product descriptions, assigned them to random product categories and associated pictures with the products.

Example of sample product data:

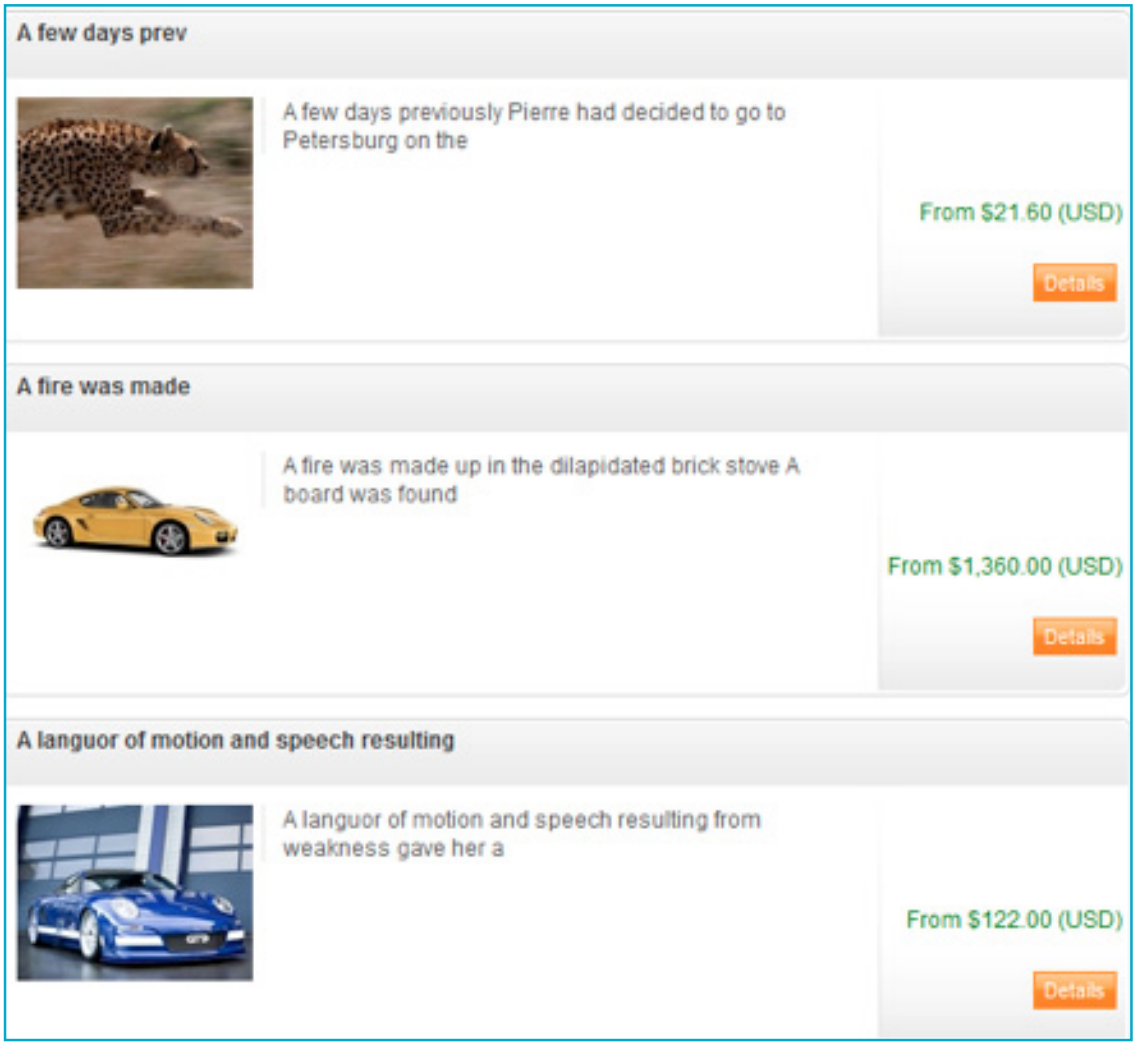

Test Platforms

We tested on four different platforms.

Physical hardware:

EUKHOSTS X3430 i5 servers with quad-core Intel Xeon (2.4GHz) processors. Each server had 8GB DDR3 RAM and local SAS drives provided storage. The servers had 64-bit Windows Server 2008 Std installed and the database was SQL Server 2008 R2 Std

VMWare:

VMware, Inc. is a company providing virtualization software founded in 1998 and based in Palo Alto, California, USA. It is majority owned by EMC Corporation. VMware's enterprise software hypervisors for servers, VMware ESX and VMware ESXi (used in this test) are bare-metal embedded hypervisors running directly on server hardware without requiring an additional underlying operating system.

Virtual DCS, our VMWare provider provisioned two virtual machines hosted on a Dell R610i Dual server with quad-core E5540 Processors. The host platform ran VMWare vSphere ESXi hypervisor with 96GB RAM, storage was provided by a HP StorageWorks P4500 G2 15k SAS SAN Array. For our test, each VM was assigned 1 x 2.25GHz core and assigned 7GB RAM. had 64-bit Windows Server 2008 Std installed and the database was SQL Server 2008 R2 Std

Other customers were sharing the VMWare platform and automatic resource allocation took place using VMWare DRS (Distributed Resource Scheduler).

Amazon AWS:

The Amazon Web Services (AWS) are a collection of remote computing services (also called web services) that together make up a cloud computing platform, offered over the Internet by Amazon.com. The most central and well-known of these services are Amazon EC2 and Amazon S3.

For our performance test, we used two "Large" Amazon instances. An Amazon "Large" instance has 4 x "Amazon Compute Units"* and 7.5GB RAM. Each instance used the Windows Server 2008 Std (64-bit) operating system and the database was SQL Server 2008 R2 Std.

*Amazon - One EC2 Compute Unit provides the equivalent CPU capacity of a 1.0-1.2 GHz 2007 Opteron or 2007 Xeon processor. http://aws.amazon.com/ec2/instance-types/

Windows Azure

The Windows Azure Platform is a Microsoft cloud platform used to build, host and scale web applications through Microsoft datacentres. Windows Azure Platform is thus classified as platform as a service and forms part of Microsoft's cloud computing strategy.

The platform consists of various on-demand services hosted in Microsoft data centres and available through through three product brands.

These are Windows Azure (an operating system providing scalable compute and storage facilities), SQL Azure (a cloud-based, scale-out version of SQL Server) and Windows Azure AppFabric (a collection of services supporting applications both in the cloud and on premise).

Microsoft has also published plans to offer the Windows Azure Platform Appliance, which can be hosted in non-Microsoft data centres. This will enable resellers, such as HP, Dell, Fujitsu and eBay, to offer cloud services based on the Microsoft Azure Platform.

For our comparison, we used an Azure "Large" instance for our web server and the database was provided by SQL Azure. *http://www.microsoft.com/windowsazure/faq/#*

Microsoft quotes a large instance as being equivalent to 4 x 1.6GHz processors with 7GB RAM

Amazon quotes a large instance as being equivalent to 4 x 1.2GHz processors with 7.5GB RAM

Monitoring the tests

We used Windows Perfmon to monitor the performance of the applications during the performance tests. PERFMON logs were then analysed using Intechnica's "KPIManager" tool.

Not all performance counters are available for Windows Azure since it operates as "PaaS", rather than "IaaS" and performance monitoring and optimisation is handled by Microsoft.

KPIManager produces charts, which like the one below, allow the comparison of key performance indicators across the target platforms.

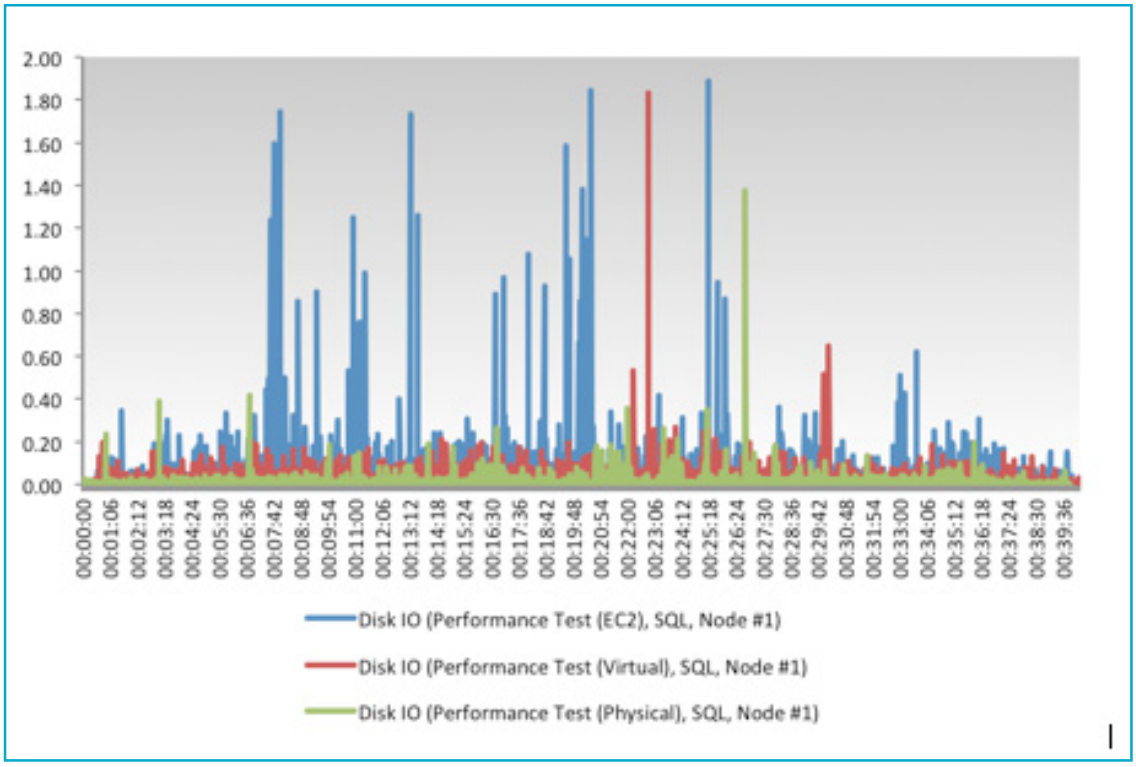

Example of KPIManager

KPI Manager also compares windows performance counters against industry norms to identify performance problems with the servers undergoing performance tests.

	Database Server KPI's	SLA RAG			Performance Test (EC2)	Performance Test (Virtual)	Performance Test (Physical)
#	KPI				SQL, Node #1	SQL, Node #1	SQL, Node #1
1	% Processor Time	>85%	50-85%	<50%	50.10	35.18	19.08
2	Processor Queue Length	>10	2-10	<2	0.93	0.57	0.16
3	Context Switches per Second	>10000	>5000	<5000	1,026.32	1,359.43	1,214.47
4	Memory Available MBytes	<100	<500	>500	6,642.57	2,422.15	6,585.86
5	Memory Page Reads per Second			<5	0.00	0.01	0.00
6	Memory Page Faults per Second				110.41	191.23	193.87
7	Memory Cache Faults per Second				0.19	0.48	0.10
8	Page File Usage %	>80%	>50%	<50%			
9	Average Disk Queue Length	>2	<2	<1	0.06	0.03	0.02
10	Access Methods\Forwarded Records/sec			< [\SQL Statistics\Batch Requests/sec] / 10	0.00	0.00	0.00
11	Access Methods\FreeSpace Scans/sec				35.22	40.18	40.44
12	Access Methods\Full Scans/sec			[\Access Methods\Index Searches/sec] / [Access Methods\Full Scans/sec] > 1000	109.89	126.59	126.53
13	Access Methods\Index Searches/sec			[Access Methods\Index Searches/sec] / [\Access Methods\Full Scans/sec] > 1000	81,769.14	91,968.05	92,127.73
14	Access Methods\Page Splits/sec			< [\SQL Statistics\Batch Requests/sec] / 5	0.20	0.23	0.23

KPIManager is used to compare windows performance counters against industry performance guidelines and identify application bottlenecks.

As well as measuring server performance, TrafficSpike together with an Intechnica application called MetricsWizard was used to compare page response times and the number of page requests served by each platform.

Example of MetricsWizard output.

Transaction	Count			Mean			Errors	
	EUKHOSTS	VDCS	EUKHOSTS<-VDCS, Variance (% diff)	EUKHOSTS	VDCS	EUKHOSTS<-VDCS, Variance (% diff)	EUKHOSTS	VDCS
<COMBINED>	9910	9915	0%	0.712	0.690	-3%	0	0
AddToCart	54	56	+4%	1.156	0.924	-20%	0	0
AddToWishlist	89	91	+2%	0.898	0.734	-18%	0	0
BrowseToProducts_Accessories	41	41	0%	0.816	0.708	-13%	0	0
BrowseToProducts_Apparel & Shoes	85	84	-1%	0.811	0.710	-12%	0	0
BrowseToProducts_Books	86	87	+1%	0.678	0.663	-2%	0	0
BrowseToProducts_Cell phones	85	85	0%	0.719	0.697	-3%	0	0
BrowseToProducts_Computers	84	84	0%	0.847	0.714	-16%	0	0
BrowseToProducts_Desktops	13	13	0%	0.607	0.641	+6%	0	0
BrowseToProducts_Digital downloads	75	75	0%	0.557	0.650	+17%	0	0
BrowseToProducts_Electronics	85	85	0%	0.724	0.641	-11%	0	0
BrowseToProducts_Games	20	20	0%	0.805	0.889	+10%	0	0
BrowseToProducts_Gift Cards	86	86	0%	0.750	0.633	-16%	0	0
BrowseToProducts_Jeans	22	22	0%	0.703	0.808	+15%	0	0
BrowseToProducts_Jewelry	80	80	0%	0.615	0.596	-3%	0	0
BrowseToProducts_Notebooks	18	17	-6%	0.570	0.743	+30%	0	0

Test results

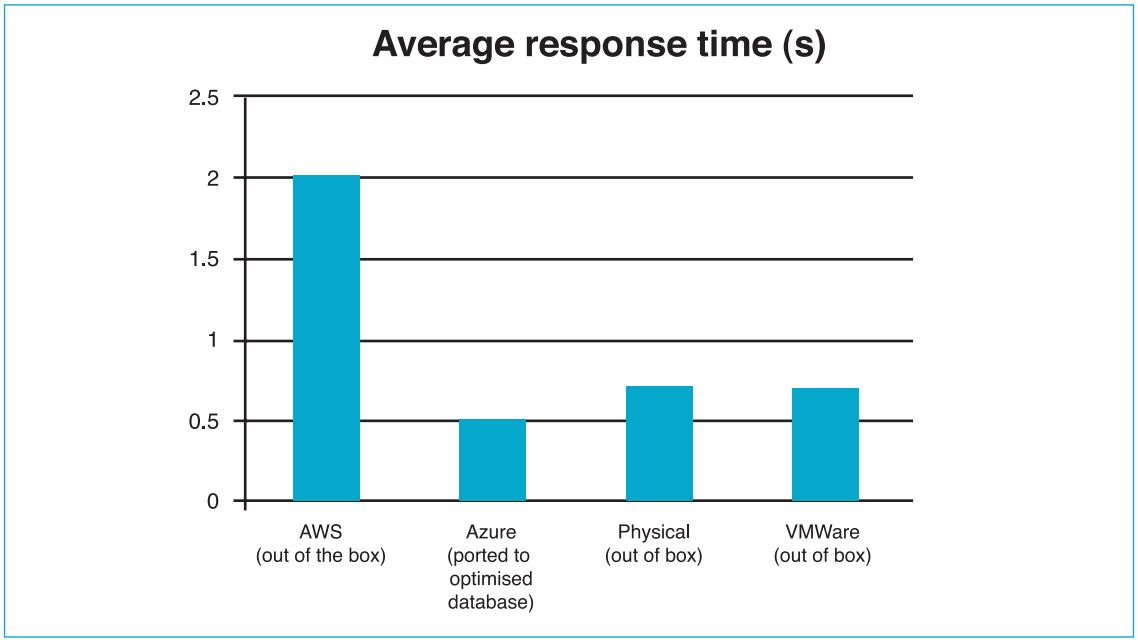

Resource utilisation

Platform	Web Server CPU utilisation* under load	SQL Server CPU utilisation* under load
AWS	57.69%	50.10%
Azure	29.73%	N/A
VMWare	27.51%	35.18%
Physical	29.79%	19%

*CPU utilisation is as a measure of total CPU available, as reported by windows performance counters.

NB Performance counters for SQL Azure not available since the database is provided as a service and performance is monitored by Microsoft. On balance VMWare has less CPU "headroom" than physical hardware based on the performance counters for the SQL server, although this is based on a single test and in reality VMWare will allocate more resources to a server at runtime (up to certain limits) if demand increases.

Other key counters such as available memory etc. were very similar across all platforms which would be expected.

KPIManager results

This example screenshot demonstrates that CPU utilisation is comparatively high on the SQL server hosted on the Amazon EC2 platform.

	Web Server KPI's	SLA RAG			Performance Test (EC2)	Performance Test (Virtual)	Performance Test (Physical)
#	KPI				IIS, Node #1	IIS, Node #1	IIS, Node #1
1	% Processor Time	>85%	50-85%	<50%	57.69	27.51	29.79
2	Processor Queue Length	>10	2-10	<2	3.72	0.81	0.31
3	Context Switches per Second	>10000	5000-10000	<5000	697.63	783.90	1,415.59
4	Memory Available MBytes	<100	100-500	>500	6,557.49	6,661.35	6,343.98
5	Memory Page Reads per Second			<5	0.29	0.59	0.46
6	Memory Page Faults per Second				990.49	382.08	683.84
7	Memory Cache Faults per Second				1.34	1.70	2.91
8	Page File Usage %	>80%	50-80%	<50%			
9	Average Disk Queue Length	>2	<2	<1	0.03	0.04	0.01
10	Queued ASP.Net Requests	>10	1-10	<1	0.52	0.11	0.09
11	Queued ASP.Net Application Requests	>10	1-10	<1	0.00	0.00	0.00

More detailed analysis of the database counters (see below) shows that the application does a large number of page lookups/sec.

This could be reduced by moving images out of the database onto the local file system and then implementing better caching on the webserver.

	Database Server KPI's	SLA RAG			Performance Test (EC2)	Performance Test (Virtual)	Performance Test (Physical)
#	KPI				SQL, Node #1	SQL, Node #1	SQL, Node #1
1	% Processor Time	>85%	50-85%	<50%	50.10	35.18	19.08
2	Processor Queue Length	>10	2-10	<2	0.93	0.57	0.16
3	Context Switches per Second	>10000	>5000	<5000	1,026.32	1,359.43	1,214.47
4	Memory Available MBytes	<100	<500	>500	6,642.57	2,422.15	6,585.86
5	Memory Page Reads per Second			<5	0.00	0.01	0.00
6	Memory Page Faults per Second				110.41	191.23	193.87
18	Buffer Manager\Buffer cache hit ratio	<90%	90%-95%	>95%	100.00	100.00	100.00
	Some rows omitted for clarity....						
23	Buffer Manager\Page lookups/sec			[Buffer Manager\Page lookups/sec] / [\SQL Statistics\Batch Requests/sec] < 100	123,860.70	139,516.20	139,724.00
24	Buffer Manager\Page reads/sec			<90	0.21	0.06	2.98
25	Buffer Manager\Page writes/sec			<90	1.07	1.34	1.00
26	Buffer Manager\Reserved pages				1,301.58	644.94	639.33

Appendix 2

Annual costs for each product offering.

	Server cost (each)	Set up cost	Server cost	Storage cost	SQL cost	Database use	Bandwidth	Admin	Annual cost
EUK Hosts	£ 189.00	£ -	£ 378.00	£ 50.00	£ 54.89	£ -	£ 399.00	£ 467.21	£ 16, 189.19
Azure	£ 439.00		£ 878.00	£ 1.34	£ -	£ 60.94	£ 450.00	£ -	£ 16, 688.16
AWS	£ 308.28	£ 1,132.04	£ 616.55	£ 47.20	£ 54.89	£ -	£ 353.05	£ 467.21	£ 19, 598.85
VDCS	£ 344.13	£ -	£ 688.25	£ -	£ 54.89	£ -	£ 500.00	£ 467.21	£ 20, 524.19

Azure	Pricing based on 2 large instances with 10GB database size (including images) Over 90% of database is product images, costs could be reduced by using file system for image storage. Data transfer costs are significant, so moving images to CDN would also save money
AWS	Pricing based on 2 large instances and 10GB database size (includes images) Could reduce database IOPS and associated charges by around 50% if images hosted externally This would also reduce data transfer costs
EUKHOSTS	Bandwidth charges are less than AWS and Azure but some price benefits would still be achieved by moving product images to CDN
Other costs	IT administrator required for patching etc. Assume 1 day/week for all except Azure, IT admin salary is £28250 Monthly cost for IT admin is £467.21 http://www.itjobswatch.co.uk/jobs/uk/it%20administrator.do

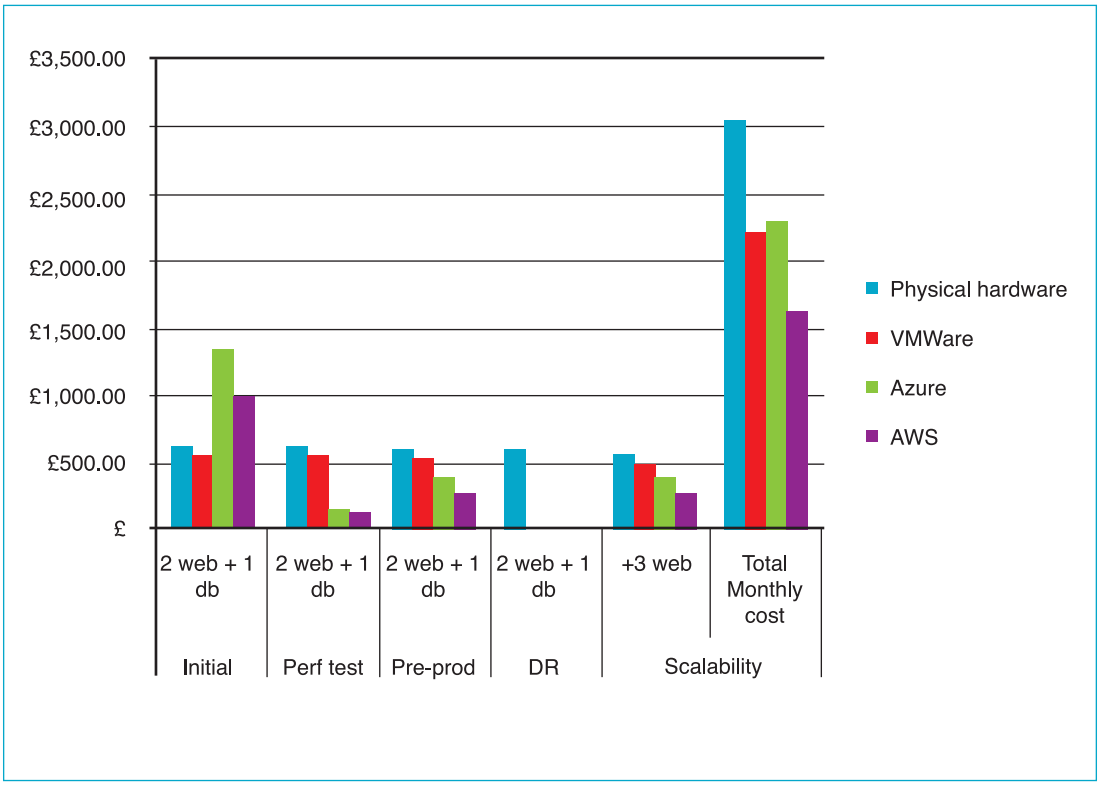

	Initial 2 web + 1db	Perf Test 2 web + 1db	Pre-pod 2 web + 1db	DR 2 web + 1db	Scalability + 3 web	Database use	Total Monthly Cost	Total Annual cost
Physical hardware	£ 622.00	£ 622.00	£ 622.00	£ 622.00	£ 567.00	£ -	£ 3, 055.00	£ 36, 660.00
VMWare	£ 571.00	£ 571.00	£ 571.00	£ -	£ 516.00	£ 60.94	£ 2, 229.00	£ 26, 748.00
Azure	£ 1, 377.00	£ 137.70	£ 413.10	£ -	£ 395.10	£ -	£ 2, 322.90	£ 27, 874.80
AWS	£ 979.00	£ 97.90	£ 293.70	£ -	£ 277.20	£ -	£ 1, 647.80	£ 19, 773.60

The table above shows the relative costs of each environment based on a performance test environment used 10% of the time and a pre-production environment used 30% of the time. Other cost savings would be likely if other development and test environments were added to these calculations.

The cost benefits of virtualisation and cloud platforms become apparent when the relative costs for these infrequently used environments are taken into account.