

Intelligent workload management

How to build, secure, manage and measure IT workloads in the age of cloud computing

December 2010

Computing infrastructure today is highly versatile and this has many benefits for businesses. Unfortunately, it also has downsides; use that infrastructure unintelligently and costs can get out of control. This is especially true as the use of both public and private cloud computing increases; it is like an all-you-can-eat restaurant – consume without thinking and you will soon become bloated.

Intelligent workload management (IWM) is the means for keeping consumption of computing resources efficient whilst still making sure business tasks are supported by all the resources they need and that this is all done securely. IWM tools support the building, management, monitoring and securing of workloads.

This report should be of interest to anyone tasked with running a contemporary computing environment who wants to be credited with placing their organisation at the leading edge of technology use within their business.

Bob Tarzey
Quocirca Ltd
Tel : +44 7900 275517
Email: bob.tarzey@quocirca.com

Clive Longbottom
Quocirca Ltd
Tel: +44 118 9483360
Email: clive.longbottom@quocirca.com

Novell.

An independent report by Quocirca Ltd.

www.quocirca.com

Commissioned by Novell

quocirca

Intelligent workload management

How to build, secure, manage and measure IT workloads in the age of cloud computing

Computing infrastructure today is highly versatile and this has many benefits for businesses. Unfortunately it also has downsides; use that infrastructure unintelligently and costs can get out of control. This is especially true as the use of both public and private cloud computing increases.

- **Intelligent management of computing is essential to support automated business tasks**
Automated business tasks comprise a number of individual workloads. These need to be built and managed intelligently and then monitored to ensure they run securely and effectively. It should be possible to change their run time environment as necessary.
- **Workloads are based around well-defined computing images**
To support a workload, an image must be built that contains all the software elements essential to it. This image can then be provisioned in a suitable environment with appropriate security. IWM is used to monitor that image at all times, allocating more resources as required and, when necessary, invoking new images.
- **What is a computing workload?**
A workload is a computing task that requires access to a mix of resources including processing power, storage, disk input/output (i/o) and network bandwidth. A workload must also understand the security and compliance needs associated with it and who has the authority to access it. To apply policies to a given workload, it must itself have a clear identity.
- **There are five basic workload types; the resources required by each varies considerably**
Application workloads run business applications and websites, database workloads handle the storage and retrieval of data, desktop workloads enable the user interface, appliance workloads deal with network and security requirements and functional workloads allow commodity tasks to be called up as web services as part of a service oriented architecture.
- **Workloads can run in three basic environments; physical, virtual and cloud**
Physical involves the deployment direct on hardware servers which are often dedicated to that one workload. With virtual, a workload is deployed onto a hypervisor that allows the easy sharing of underlying physical resources. Cloud is a massively virtualised platform, either provided by third parties or run internally, which may have thousands of underlying server, storage and network assets.
- **Better security and compliance are fundamental benefits of IWM**
This is not just about the securing of workloads as they are built and making sure that they remain secure once deployed. It is also about understanding the compliance issues of storing data in the cloud and recognising the potential dangers introduced by user-initiated workloads such as on-demand office tools.
- **The cost of deploying IWM must be offset by the benefits to the business**
The total value proposition for IWM includes reducing the cost of IT infrastructure usage, reducing the risk of running automated business tasks and the creation of incremental value. All three issues need to be considered and offset against the cost of deployment when building the case for investment in IWM tools.

Conclusions

There will be few organisations without automated tools for IWM in place that will not recognise many of the benefits of deploying them. As the world moves to an environment where more and more computing resources are available as on-demand commodities, the need to use them intelligently will only increase.

1. Introduction – the changing IT workload

In the early 1990s it became possible to walk into an office and observe that a business task that was supported by computing was being carried out on multiple machines. Users were starting to have their own desktop computing devices that drove their interface and somewhere behind these scenes were larger computers running shared applications that did the backend stuff. Client-server computing had been born; many computer-driven tasks had been split into two distinct workloads.

In those days computing was largely an internal affair; different computers were linked together by local area networks and different premises of a company were linked together by private connections. There was little connectivity to the outside world, apart from the odd dial up-service, and therefore little need for network security aimed at keeping out malevolent outsiders.

Move forward twenty years and the way computing environments are configured and run has evolved beyond recognition. Users can be granted access to applications from almost anywhere on a bewildering range of devices. The central applications will often be running on virtual platforms consisting of multiple physical servers and storage arrays that may be owned, housed and run by third parties and access is often over shared public network infrastructure. Anyone doubting the increase in the number of physical devices need only look at the figures for an example end user organisation (Figure 1).

This change has happened because the technology can support it and businesses can benefit from the flexibility offered by such computing environments at a lower and lower cost per unit of compute power.

Many business tasks have become highly automated and the individuals from the multiple organisations involved in them are highly connected. This drives up business efficiency and lowers the cost of doing business. Or at least it should do.

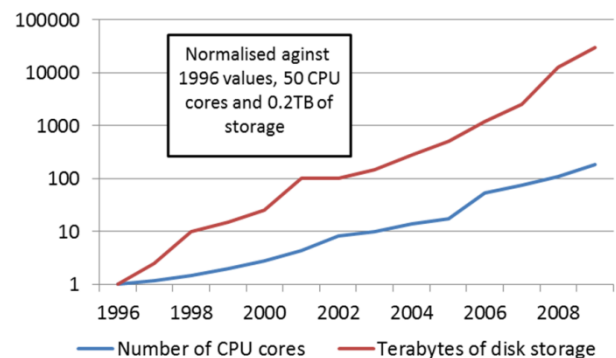
There are two major downsides. First, if IT infrastructure is not well managed, the cost benefits of a flexible platform evaporate through waste and confusion. Second, if the environment is not secure, the use of open networks means the infrastructure is vulnerable to attack from outsiders set on stealing personal information or intellectual property.

To avoid these pitfalls it is necessary to ensure that the different workloads that support given business tasks are running in an environment that best suits them and that the environment is secure and only accessible by authorised users.

For example, the resources needed to run a database and a user desktop are very different, yet, with the growing use of virtual desktops, both could end up running on the same physical infrastructure in a virtualised environment that is not intelligently managed.

This report aims to explain how computing workloads can be built, managed and monitored intelligently to ensure that businesses get the flexibility and cost benefits expected from 21st century computing, without compromising security. It should be of interest to anyone tasked with running such environments who wants to be credited with placing their organisation at the leading edge of technology use within their business.

Figure 1: Relative increase in CPU and storage since 1996, European Bioinformatics Institute¹




“If IT infrastructure is not well managed the cost benefits of a flexible platform evaporate through waste and confusion”

2. What is a workload?

A workload is a computing task that requires access to a mix of resources including processing power, storage, disk input/output (i/o) and network bandwidth. The workload can be purely technical (e.g. supporting the operating system or running anti-virus software) or be business focused (e.g. enabling customers to order on an ecommerce site or end users to create documents). A workload must also understand the security and compliance needs associated with it and who has the authority to access it. To apply policies to a given workload it must itself have a clear identity.

There is nothing new about defining IT in terms of workloads; it is just that the flexibility with which workloads can be deployed and moved has increased massively in the last ten years or so. This has come about because of two revolutions in the way computing environments are managed; virtualisation and cloud computing. This means there are now three basic environments in which a work load can be run:

- Physical: workloads are provisioned directly onto hardware in the form of dedicated servers.
- Virtual: workloads are provisioned on a hypervisor that hides underlying hardware resources from the workload. This allows different workloads to easily share physical resources, providing better resource utilisation and economies of scale.
- Cloud: workloads are provisioned on massively virtualised platforms where tens, hundreds or thousands of servers are pooled together. Clouds were initially largely provided as a commodity resource by third parties, but are now also increasingly being seen as a good way to manage privately owned computer infrastructure.



“Ensuring that a given workload runs in the environment that most suits it at any given time is the aim of intelligent workload management”

A workload or a collection of workloads makes up a business service, which is what the end user consumes. For example, a business service, such as a customer relationship management (CRM) application, could consist of a database workload on a physical server plus an application server workload on a private cloud plus a presentation layer workload in the public cloud.

There are five basic types of workload:

- Application workloads: such as a CRM application or a website.
- Database workloads: the management of backend data storage such as a SQL database or enterprise content management systems.
- “Desktop” workloads: the term is used loosely, and taken to mean any user access environment. This could be a PC running Windows or a smaller form factor device such as a netbook or smartphone. Increasingly, it also includes virtual desktop infrastructure (VDI). See section 4 for further information.
- Appliance workloads: many security and network functions, for example firewalls, have been carried out in the past using dedicated hardware appliances. Many vendors have stopped providing the hardware themselves and now supply virtual appliances. It is then down to the user organisation to decide how best to deploy them – on dedicated hardware or as virtual machines.
- Functional workloads: web services provided by third parties, or internally, that are called up over the network and utilised by other workloads. This is the basis of service oriented architectures. Such services can also become an issue if invoked directly by users without the knowledge of IT management.

Ensuring that a given workload runs in the environment that most suits it at any given time is the aim of intelligent workload management.

3. The role of intelligent workload management (IWM)

IWM provides the capability to build, secure, manage and measure images.

Building and securing images requires selecting the required software components (which may include an operating system, application server, application, database and security software) and then building static images or enabling them to be created on the fly. Upgrades and patches need to be regularly applied to the software components that comprise an image and the dependencies between different software layers understood and managed (for example what operating system version supports a given application version?). It is also necessary to ensure a given image meets the security and compliance needs associated with it and that the software licences are valid, up to date and used efficiently.

Having built an image, IWM enables its management. First the initial run time environment for the image must be defined and allocated. Second, its requirements need monitoring through time and the environment changing as necessary or, on occasions, the image moving to a different environment. IWM allows the effects of such changes to be modelled in advance to ensure that when changes are made they are effective. When the time comes and the image is no longer need, it needs to be safely deprovisioned.

IWM tools allow all this activity to be measured and reported on for both technical and business staff. For the former, this helps ensure agreed service levels are met and that the use of data is secure. For the latter, it spells out the cost to the business of running the various applications that support it and enables compliance reporting.

It is often assumed that virtual and cloud environments make ensuring that an image has the resources it needs easy, but that is not always so. It is true that allocating more processing power and storage is usually straight forward; however, ensuring there is enough i/o and bandwidth is harder to predict and manage, and may involve invoking multiple images for a given workload and then balancing use across those images or moving it from a virtual to a physical server.


4. The evolution of workload management

There is nothing new about the use of virtual computing platforms; indeed the term is imbedded in computing history. Mainframe computers have been running IBM's MVS (Multiple Virtual Storage) for years. In the 1980s DEC introduced VAX (virtual addressing extension) mini-computers running VMS (virtual memory system). Such devices, especially mainframes, are still running computing workloads and for many will be a resource for IWM.

Virtualisation has always been required to allow the sharing of IT assets. The difference today is that virtualisation is being used to virtualise commodity x86 servers and, in the process, allow the construction of huge cloud computing platforms. Provision a workload on a mainframe and you can walk into the data centre and touch the device on which it is running. Provision that same workload into a virtual cloud environment and you would have no idea on which of the physical servers underlying the platform the workload was actually running on.

Virtualisation has been used by managed hosting and other cloud platform providers to build public computing platforms. This allows organisations to rent compute power, in addition to that which they own in house, as a commodity service, often accessed over the public internet. These public clouds are, by definition, shared by multiple different organisations. This makes the need to make sure a workload has the required resources and is running securely even more acute. Figure 2 shows how the use of public cloud is expected to increase over the next few years.

Figure 2: Growth in use of cloud – source Nov 2009, Goldman Sachs IT Spending Survey²




Public cloud services fall into one of three basic categories:

- The lowest level is infrastructure-as-a-service (IaaS). Workloads are provisioned at the hypervisor level, i.e. the workload must have its own operating system. IaaS is particularly useful for organisations that are running virtualised applications internally, but may want to make use of additional capacity when their own resources are stretched (examples are Amazon EC2, Rackspace Cloud Servers, Savvis Symphony and Attenda RTI).
- Platform as a service (PaaS) goes a stage further and includes the operating system and other application services. PaaS suits organisations that are committed to a given development environment for a given application but like the idea of someone else maintaining the deployment platform for them (examples include Microsoft Azure, Google App Engine and salesforce.com's VMforce). Most managed hosting providers also provide PaaS-type offerings for both Windows and Linux environments, although these will often not be fully shared infrastructure (see Quocirca report "Managed hosting in Europe"³)
- Finally, there is software as a service (SaaS). This goes the whole hog, offering fully functional applications on-demand to provide specific services such as email management, CRM, ERP, web conferencing and an increasingly wide range of other applications. SaaS may seem to be beyond the scope of IWM as there is no capability to choose how an application is deployed. However, from the point of view of monitoring, it is important to understand what SaaS workloads users are invoking for themselves. For example, if a user invokes Google Apps to create a document which includes personal information, where does it end up being stored and might this be in breach of data protection regulations? Individuals create data but enterprises are responsible for security, privacy, reliability, and compliance.

There are obvious benefits to being able to buy computing resources as and when they are needed and making sure such resources are procured and used judiciously is part of IWM. The effectiveness of the technology developed to deliver public cloud computing is being recognised by IT departments to the extent that many of them are now turning their own virtual computing environment into private clouds. This reduces the burden of day to day management even more, but drives up the need for IWM.

Some organisations are now starting to combine private and public clouds to create hybrid clouds (Figure 3). This allows them to maximise the use of resources they own first and then to spill over and procure public cloud resources only when the overall workload requires it. There are other benefits; hybrid cloud is a useful way of providing additional resources whilst new private infrastructure is being deployed. It can also be used as a step in the migration from private to public cloud; new functions are brought in from the public cloud, existing internal functions are reviewed on a regular basis and considered for migration to, or replacement with, a public cloud function. However such functions are provisioned they are still workloads that need monitoring.

Figure 3: Public clouds, private clouds and hybrid clouds



One of the potential downsides of cloud platforms is server-sprawl. This can be likened to an all-you-can eat restaurant menu. If the way resources are consumed is not carefully managed, usage will soon become bloated. This is in the interests of public cloud providers; more usage means more revenue, whatever the reason. However, the revenues received have to offset the incremental costs of hardware acquisition, management and replacement. If the provider's costs are not competitive in the market, then the customer may well take the opportunity to move to a better managed environment where the costs are lower.

With private cloud, usage and management costs are driven up and what should be an efficient way of managing hardware resources does not live up to expectations. IWM helps mitigate server sprawl in both public and private cloud environments.

There is one type of workload that is particularly problematic, the “desktop”. The term is used loosely to define the means by which a given user interfaces with computer-driven tasks. The range of user environments is now so broad compared to the hegemony of the Microsoft desktop just 10 years ago that management and security has become a big issue. This is leading many to bring the desktop workload back to the core and just use the device as a way of accessing it rather than actually running it.

Managing so-called virtual desktop infrastructures (VDI) will be an increasingly important part of IWM going forwards; analyst group Gartner has predicted the number of virtual desktop units will rise from 500,000 in 2009 to 49 million in 2013³. Indeed, the need for virtualised desktops to be flexible in how they provide solutions to different users means that IWM is required to look at what runs where – for example, is everything suited to be run at the server, should some of it be provisioned to the desktop, or should the image be centrally controlled, but run completely at the client?

VDI also helps support another growing trend in corporate computing; consumerisation. This is the desire for employees to use their own devices to access corporate computing resources. This, in part, helps explain the growing complexity of the range of end-user devices IT departments have to support. Managing and securing an individual user’s own device is challenging, so it is easier to provide secure access to a virtual desktop that is managed centrally.

“The number of virtual desktop units will rise from 500,000 in 2009 to 49 million in 2013”

IWM also helps enable a key element of business continuity. As business tasks have become more and more reliant on compute power it is necessary to ensure that most compute workloads are continuously available or at least can be redeployed very quickly. For example, should a database workload running on a physical server fail, IWM tools could allow it to be redeployed on virtual infrastructure whilst the problem is fixed; the business task continues, even if there is a short term impact on performance. Of course, many will keep a backup database server on “hot-standby”, fully functional and ready to go. However, IWM also helps enable “warm-standby”, where a backup workload is ready to go in a virtual environment, separate to that of the primary workload, but is not consuming full resources until required to do so.

5. Ensuring workload resources

The environment in which a workload is deployed will depend on the physical resources it is likely to require. These requirements will also vary through time and it is necessary to monitor workloads constantly to ensure they have the resources they need and to take the necessary actions to provide extra resources. This may mean:

- Allocating more resources
- Starting new instances of the workload
- Moving the work load to a different environment, e.g. from a virtual environment to a dedicated physical server.

There are four basic physical resources requirements that a workload must draw on; compute power, i/o, network bandwidth and storage. These are characterised by examples below:

- Compute power: Some applications are more reliant on compute power than anything else, for example the number crunching applications used by oil companies to understand geological structures or telephone companies wanting to analyse user behaviour. There are times when such applications have peak requirements over and above that available internally and this is a classic use case for on-demand compute power provided from the cloud.
- i/o: a database is an i/o intensive application that must have sufficient resources to serve hundreds or thousands of concurrent users very quickly. It should not have to compete for resources with virtual desktops writing copies of bulky documents to disk every five minutes. To this end, high performance database workloads are still often best run in dedicated physical environments where resources are guaranteed.

- Network bandwidth: a video conferencing server needs a lot a bandwidth to ensure the user experience is of the highest quality. It needs that bandwidth in real time and should not be forced to queue with an email server that may be dealing with non-urgent attachments to messages. It may make sense that the video workload is also provisioned in a dedicated physical environment where priority bandwidth can be guaranteed.
- Storage: a VDI environment running hundreds of desktops will need a lot of storage allocated. It is no good if a user has just spent hours preparing a customer presentation only to find there is no space to store it. Thin provisioning of storage avoids having to pre-allocate too much disk space; IWM enables more to be added just ahead of the predicted requirement. It also makes sense to minimise the data storage associated with workloads, especially when they are running in public clouds. Only capture and hold information essential for the workload to operate.

This need to understand the requirements of workloads and select the best environment for them to run in where appropriate levels of physical resources levels can be guaranteed is achievable with IWM; providing both the flexibility and control needed to maximise the use of resources and ensure service levels.

IWM also allows workloads to be moved around from one type of environment to another. It may make sense if a holiday peak is coming to move a website to a dedicated physical server where bandwidth can be guaranteed; it may be that at that time of year a normally busy database has a reduced number of transactions. Swapping the two around for a period of time may make sense. The ability to move workloads between physical and virtual environments is a key benefit of IWM.

The flexibility to move workloads is all well and good, but there is some additional intelligence that needs to be applied before doing so. The availability of public cloud resources is a boon for those needing to manage unpredictable demand, but the provenance of the public cloud needs to be understood to ensure compliance with data privacy requirements. For example, the UK Data Protection Act states that personally identifiable data should not be held outside of the EU. Some public cloud providers do not offer such guarantees and, even if they appear to, there may still be problems when you look behind the scenes at their redundancy and data backup processes. One way around this is to use role mapping to identify a user's personal data to customise their experiences, without exposing the actual personal data to the cloud.

“The provenance of the public cloud needs to be understood to ensure compliance with data privacy requirements”

Consideration also needs to be given to what happens to the footprint of a workload once it is de-provisioned from a public cloud environment. Is sensitive data being left behind? If a cloud provider does not guarantee that this will not happen, it may make sense to encrypt data in the first place, rendering it useless to those without the keys. Many may feel more comfortable with data stored in a public cloud environment being encrypted regardless of any guarantees offered by the cloud provider.

6. Conclusion – a total value proposition for IWM

Deploying IWM has a cost. This cost needs to be outweighed by the benefits to the business; some of that will come through reduced costs of running actual workloads, but there are other, less tangible ways, that IWM will create value for an organisation and, of course, as already discussed, IWM has a role in ensuring security and therefore decreasing risk. Quocirca uses a total value proposition model to understand how these factors play out for an IWM deployment.

All these factors will vary from one organisation to the next, but there will be few organisations without automated tools for IWM in place that will not recognise many of the above benefits. As the world moves more and more to an environment where computing resources are increasingly available as on-demand commodities, the need to use them intelligently will only increase.

Considerations for building a total value proposition for an IWM investment

Cost of deployment of IWM

- Cost of tools
- Cost of platform to run tools on – backend server probably negligible, the tools themselves are likely to free up far more resources than they use
- Cost of deployment – consultancy, training etc.

Cost saved by deploying IWM

- Better resource usage – therefore less need to throw money at the problem through investment in physical assets
- Fewer physical assets requires fewer system management staff
- Better resource usage means less power consumption
- Intelligent use of cloud resources – only invoking them when needed
- Reduced cost of software licencing through better management; needlessly invoking workloads, or invoking them in an over-specified environment can lead to increased licence charges from some vendors
- Automation of mundane IT tasks through the use of IWM frees up IT staff for other tasks and decreases the overall staffing requirement

Value created through use of IWM

- Reduced power consumption and better resource usage provides a feed into sustainability reporting
- Reduced IT outage increases productivity
- Improved support for consumerisation; the desire for employees to use their own devices to access corporate IT
- Faster response to technical needs through faster provisioning and better prediction of workload needs
- Support for migration from physical to virtual environments and from internal and external clouds

Risk reduction through use of IWM

- No loss of workloads through resource failure
- Security of workloads more easily monitored and guaranteed
- User-initiated workloads can be monitored and any risk they pose to sensitive data monitored and controlled
- Data breach avoidance
- Capability to carry out “what if” scenarios without impacting the run time environment
- Monitoring of user initiated workloads (e.g. Google Apps)

7. References

1. Figures published by SNS Europe, Volume 10 Issue 4, Summer 2010 – Interview with Head of IT at the European Bioinformatics Institute
2. Managed Hosting in Europe, Quocirca <http://www.quocirca.com/reports/16/managed-hosting-in-europe--june-2009>
3. Goldman Sachs Global Investment Research; A paradigm shift for IT: The Cloud – November 2009
4. www.connectitnews.com; Global hosted virtual desktop market to surpass \$65 billion in 2013, 12 April, 2009 <http://www.connectitnews.com/usa/story.cfm?item=3173>

For more information on intelligent workload management go to www.intelligentworkloadmanagement.com where a regular free newsletter is available.

About Novell

Novell, a leader in Intelligent Workload Management, helps organisations securely deliver and manage computing services across physical, virtual and cloud computing environments. We help customers reduce the cost, complexity, and risk associated with their IT systems through our solutions for identity and security, systems management, collaboration and Linux-based operating platforms. With our infrastructure software and ecosystem of partnerships, Novell integrates mixed IT environments, allowing people and technology to work as one.

Novell's approach to IWM is called WorkloadIQ which helps:

- Create a flexible and agile IT infrastructure.
- Leverage the cost savings of virtualisation and cloud computing.
- Maintain control and security across all your computing environments.
- Ensure a seamless computing experience for your end users.

Four Key Lifecycle Phases; WorkloadIQ enables you to build, secure, manage and measure workloads across physical, virtual and cloud environments.

- **Build:** Build intelligent workloads that are portable and have identity, security and management services integrated with the application.
- **Secure:** Ensure the right levels of data protection and regulatory compliance while controlling access across multiple computing environments.
- **Manage:** Manage and move workloads within and across all computing environments to optimise workload execution and utilisation of IT assets.
- **Measure:** Measure workload performance and monitor security events to generate a real-time, identity-aware view of your IT infrastructure.

A Modular Approach at Your Pace.

Concerned about the potential cost or complexity? Don't be.

Novell provides a full range of WorkloadIQ solutions and products. You can get started with just one or two WorkloadIQ offerings while we help you leverage your existing technology assets. Over time, you can move toward the complete WorkloadIQ vision at a pace that makes sense for your IT organisation.

Your enterprise needs both flexibility to meet today's changing business demands and better control of computing resources across physical, virtual and cloud environments. Novell—along with our broad ecosystem of partners—can get you there.

For more information about Novell Intelligent workload management go to www.novell.com/workloadiq

The Novell logo is displayed in a bold, red, sans-serif font.

REPORT NOTE:

This report has been written independently by Quocirca Ltd to provide an overview of the issues facing organisations seeking to maximise the effectiveness of today's dynamic workforce.

The report draws on Quocirca's extensive knowledge of the technology and business arenas, and provides advice on the approach that organisations should take to create a more effective and efficient environment for future growth.

Quocirca would like to thank Novell for its sponsorship of this report.

About Quocirca

Quocirca is a primary research and analysis company specialising in the business impact of information technology and communications (ITC). With world-wide, native language reach, Quocirca provides in-depth insights into the views of buyers and influencers in large, mid-sized and small organisations. Its analyst team is made up of real-world practitioners with first-hand experience of ITC delivery who continuously research and track the industry and its real usage in the markets.

Through researching perceptions, Quocirca uncovers the real hurdles to technology adoption – the personal and political aspects of an organisation's environment and the pressures of the need for demonstrable business value in any implementation. This capability to uncover and report back on the end-user perceptions in the market enables Quocirca to advise on the realities of technology adoption, not the promises.

Quocirca research is always pragmatic, business orientated and conducted in the context of the bigger picture. ITC has the ability to transform businesses and the processes that drive them, but often fails to do so. Quocirca's mission is to help organisations improve their success rate in process enablement through better levels of understanding and the adoption of the correct technologies at the correct time.

Quocirca has a pro-active primary research programme, regularly surveying users, purchasers and resellers of ITC products and services on emerging, evolving and maturing technologies. Over time, Quocirca has built a picture of long term investment trends, providing invaluable information for the whole of the ITC community.

Quocirca works with global and local providers of ITC products and services to help them deliver on the promise that ITC holds for business. Quocirca's clients include Oracle, Microsoft, IBM, O2, T-Mobile, HP, Xerox, EMC, Symantec and Cisco, along with other large and medium sized vendors, service providers and more specialist firms.

Details of Quocirca's work and the services it offers can be found at <http://www.quocirca.com>