BUSINESS INSIGHTS July 2007

Contacts:

Clive Longbottom
Quocirca Ltd
Tel +44 118 948 3360
Clive.longbottom@quocirca.com

Dennis Szubert
Quocirca Ltd
Tel +44 845 397 3660
Dennis.szubert@quocirca.com

Data Centre Automation

Enabling the Agile Enterprise

Executive Summary

From humble roots in basic provisioning, data centre automation has grown in response to years of unbridled server proliferation, and is now transforming the enterprise data centre

Main Findings

- **Data Centre Automation (DCA) is becoming a necessity**
  The increasing complexity of the data centre and dependent systems means that tooling is now required to keep track of, and maintain the environment in a manner which involves as little human intervention as possible
- **Freeing up expensive human resource is becoming all important**
  With upwards of 70% of an organisation's IT expenditure being spent on maintenance and fire fighting, tooling is required to drive down costs. The availability and cost of adequately skilled human resources is becoming prohibitive – and utilising these skills for maintenance is not cost-effective. Freeing up these resources for value-add strategic IT implementation work is a necessity
- **Discovery, Dependency mapping, Maintenance and Dynamic Provisioning are all key areas**
  Many organisations have no real knowledge of how many servers and other hardware resources they have in their environment – never mind to what level of version and patch each layer of the stack may be at. DCA tooling can automate each of these tasks, and ensure that systems are maintained at necessary critical patch levels and optimum operating levels
- **Utility computing will drive the need for greater flexibility in the data centre**
  Service Oriented Architecture (SOA), Software as a Service (SaaS) and other utility concepts such as grid computing will increase the need for a fully dynamically managed infrastructure. Manually managing and maintaining such environments should not be considered – DCA provides the best way to create a managed system.

Conclusion

The opposing forces of greater demands from the business for increased flexibility and speed of response compared with the drive for lower infrastructure costs while complexity increases means that IT departments have to look to solutions to replace the costly and generally inefficient human aspect of managing the data centre and its dependent systems. DCA provides the way to create a highly dynamic and responsive capability that provides increased support for the business.

An independent report by Quocirca Ltd.

www.quocirca.com

Commissioned by Symantec

1 Background

Organisational computing and data needs have moved in a cyclical basis – in the age of the mainframe, everything to do with the computing environment was held within one environment, the data centre. Then there was the move to distributed computing, with departmental computers and data being held more locally to the users. However, as data volumes grew and the organisation became more dependent on the IT infrastructure, issues caused by lack of easy access to disparate data sources meant that the capability to deal with data or infrastructure loss was impaired. This led to a move back to monolithic data centres, with massive data warehouses holding aggregated data alongside racks of servers dedicated to specific tasks. However, different enterprise applications required different data stores, and silos of data began to appear within these data centres, and the speed of response when interrogating and reporting against these data stores became such an issue that departments started to again look to localised computing as a means of supporting their own needs.

Now, the move is towards a service oriented architecture (SOA), complete with data federation (the capacity to bring multiple different data sources together for reporting and business intelligence purposes) and master data models (the capacity to have a single view of the truth for common data through a single database for information such as customer records). The increase in interoperability due to greater standards adoption within the vendor community, combined with the growth in understanding and adoption of technologies such as server and storage virtualisation is making such moves far more cost effective and attractive than has previously been the case.

It is unlikely that we will see the data centre of old reappear – although many organisations will have a large centralised environment where a greater part of their computing environment will reside, there will always be data and servers that are held external to this environment – and yet this data still has to be included in any data management and reporting approach. To this end, organisations will need to invest in the tools required to ensure that a highly flexible data and infrastructural management environment is created for the future.

This document looks at the needs of the organisation around existing and future computing and data needs, and at possible approaches to ensure that these needs are met in a flexible and effective manner.

2 Business Issues

Today's business environment has changed rapidly, even compared to just a few years ago. Whereas a high street retailer would know who their competitors were – mainly by looking up and down the high street itself – today, these same companies are continually having to watch over their shoulders at what is happening on the internet, as they see short-lived, but low-cost competitors come and go. With the internet offering the capability for new business models to be tested at low capital and resource cost compared to the standard costs for existing organisations with large legacy infrastructures, the business impact of these highly leveraged, low resourced "global" companies cannot be underestimated.

More than ever, enterprises are looking to reduce operational costs, improve quality, speed execution times, and increase security. However, these challenges can directly impact the ability to respond to business needs and the latter areas are often at odds with the driver to reduce operational costs. Many IT organisations are only able to keep up with today's operational demands by continually increasing staff numbers and infrastructure assets (e.g. servers, switches, storage), and by buying significant amounts of additional tooling. Not only is that strategy costly, but it is also unwieldy - large scale, business critical changes can often take weeks or months to implement.

Data volumes are also increasing, due not only to organic growth in transactions, but also through the increasing burden of governance and compliance. From a business point of view, the simplistic approach of aging data through levels of hierarchical storage management (HSM) based on date of creation/last use and then deleting "old" data is no longer a viable option. We must be able to start from a position of understanding the lifecycles of the data,

the legal requirements for maintaining that data, the business implications and needs, and then build policies and procedures that make the greatest sense for the business.

Businesses are also facing the need to be able to react far more rapidly to change – workloads against existing IT assets can change far more dramatically than in the past as, for example, new marketing campaigns drive prospects to the web site or to the contact centre. The need to have a flexible IT environment that enables such workloads to be easily managed is becoming a survival factor – those who cannot meet business needs with technical facilitation will fail.

Today's organisations have to be flexible and fleet of foot just to survive. The capability to rapidly provision and reprovision the infrastructure to facilitate the changes in an organisation's business processes is essential. That this provisioning and managing of the data centre environment has to be carried out at the lowest cost is a given – and therefore, organisations are looking at as much automation as possible, leaving out the expensive human interventions apart from where exceptional circumstances dictate.

3 Data Centre Pain Points

Every enterprise data centre faces the challenges of trying to meet their business' demands - provisioning servers for new applications, patching hundreds of servers to protect against security breaches, or making frequent changes to multi-tier applications. With the explosion of new technologies, server complexity and continued heterogeneity, these responsibilities have become more complex and demanding. The proliferation of new software, the introduction of more complex applications and application architectures such as SOA, combined with the rate of application changes have put new pressures on IT groups within organisations.

In addition to battling to meet the demands of the business - managing constant change and meeting rising quality expectations - there are three critical issues that rise to the top of most organisation's priority lists:

- **Complexity**
  After years of unbridled server proliferation, with a separate server being added for each additional application - accounting, supply-chain management, CRM, or human resources - the net result is the data centre is increasing in complexity daily. Most of today's IT environments contain multiple operating systems, applications and hardware configurations, have far too many people and technology touch-points, and - most importantly – they have numerous integration requirements. The challenges that IT professionals face in provisioning, operating and updating such large, diverse application environments are numerous, and error rates are high. Managing configuration change across such a mix of systems can be a nightmare. It requires a knowledge of all the servers and applications in the environment, as well as the complex dependencies among them; consequently the risk of application downtime or poor performance resulting from inaccurate configuration changes are high.

- **Costs**
  As a direct result of this complexity, Quocirca's research shows that for organisations where the business and the IT group are only loosely coupled, up to 75% of the costs of operating a corporate computing system can come from maintenance and fire fighting.

  The manual monitoring and intervention for the support of both servers and applications is inefficient and expensive. Further, the "server dedicated to a single application" approach has not only led to complexity, but also to a vastly underutilised computer resource. Again, Quocirca's research has shown that most single application servers are used at less than 20% of total capacity, averaging around 10-15%. That's obviously wasteful – not only in the underutilisation of possible hardware resources, but in the need for power, for cooling, for real estate to house the hardware, for operating system and other base platform licensing and maintenance fees as well as the number of administrators and support staff required to keep these systems operating.

- **Delays**
  Installing and configuring packaged software from the hardware up – operating systems, database servers, messaging servers, web servers, application servers, and other software infrastructures — is one of the most common and time consuming tasks system administrators perform. After basic installation of packaged software, administrators can spend days or weeks configuring the software components to make systems 'work' in the manner required. Patching and upgrading these systems is also time consuming – and can be fraught with problems as support staff identify problems with disk space, device drivers and application compatibility.

Computing can, and must, be made less complex and more efficient. Companies need help to manage all that complexity, and they are increasingly looking to automation software to provide that help. IT departments are faced with pressure to improve service levels while decreasing operational costs; data centre automation provides a perfect solution for achieving such levels of operational efficiency.

4 What is Data Centre Automation?

Data centre automation (DCA) is the process of taking the ad hoc, manual tasks that typically consume the majority of the time of IT staff and transform them into automated, systematic processes managed in a more consistent, predictable manner than previously. Software is used to automate the manual tasks required for day-to-day operations of the data centre, such as: -

- **User provisioning**
  User provisioning is the addition and management of user accounts and the provision of capabilities to access certain functions, applications and data securely. Also, user provisioning covers the management of changes to a user's account, such as change of passwords, names or role within the company, etc.
- **Software & Operating System Provisioning**
  The building of full technical function capability onto a "bare metal" environment (a basic server) or on to a "warm" server (one with the base level stack of operating system and sometimes application server already present)
- **Application Deployment**
  The controlled roll out of new applications or functions to existing environments to meet the needs of the organisation
- **Patch Deployment**
  The controlled rollout of changes to existing applications or functions on a targeted or broadscale level
- **Change Management**
  The orchestration of all the changes above, ensuring that any possible impact is known prior to any change taking place, and that systems can be easily rolled back to a known position should any problems arise during the process of the change
- **Application Availability**
  Once applications have been installed, availability has to be monitored and maintained. This may be through reprovisioning an application from a failed physical hardware asset to another part of the data centre, or more often will be through the automated fixing of the application through patching corrupted files via the likes of patch deployment.

Data centre automation grew out of "server automation" - simple provisioning tools - into general-purpose platforms for configuration management and automation of IT operations. The advent of clustering and virtualisation have further transformed DCA, taking it beyond automating system changes in a relatively static environment to fit into the trend for utility or on-demand computing. Applications and their related data can be routed to different servers, storage devices, or sections of a network depending on where excess capacity exists at any moment, with minimal human intervention, maximising both server utilisation and application availability. Service-based components, commonly utilised within a service oriented

architecture (SOA) can be dynamically provisioned to support changing workload requirements driven by market and campaign needs.

5 Why Automate?

The typical way to manage an IT infrastructure is to build a team of experts - system administrators, network engineers, database administrators, operations engineers, etc. - to manage day-to-day operations. As the complexity of an organisation's IT infrastructure increases, the operational expertise required to maintain it also increases. Skilled human resource in many IT areas is a scarce commodity, and ensuring that sufficient skills are available is a costly business. For many, the obtaining and maintaining of such skills is seen as too high a cost, and the skills available to them are therefore often insufficient to meet the business needs. By using automation software to automate common, repeatable tasks, enterprises can achieve unprecedented results in the areas of efficiency and productivity, cost reduction, speed of deployment, problem diagnostics and risk mitigation. This then frees up the expensive human resource to carry-out more strategic IT work focused on facilitating the changing business process needs, rather than focusing on maintenance and fire fighting.

The following areas demonstrate where and why automation can have the biggest impacts for a business:

- **User Provisioning**
  Adding users to a company, and then managing the changes has historically been a time and cost consuming issue. However, by automating this task and making changes a user self-service task, cost savings can be massive. Further, end user self service automation tends to create greater user buy-in to an environment, as they perceive that they have greater control over their own environment – and that changes can be made far more rapidly than when having to go through a help desk or other centralised facility.

- **Software Provisioning**
  Automation software provides a consistent way to install, configure, update and remove packaged software across servers in the data centre, enabling one administrator to share lessons learned and best practices across the entire organisation. This can mean that no one in your organisation will have to reinvent the wheel for each new application deployment, preventing the ramifications of unsuccessful installations and poor configurations and the resulting application instability and downtime. "Golden" images can be created that combine all the components of a solution stack – OS, application server and application/function - making provisioning a fast and effective process with guaranteed results.

- **Configuration Management**
  Code and configuration changes, whether driven by internal changes or through a hardware vendor's update notifications, are a frequent cause of application downtime or poor performance in an organisation's operations infrastructure. The need to distribute configurations to applications occurring across multiple operating systems, including Windows, Linux and various flavours of proprietary Unix creates an even more complex task for administrators. Automation systems will be able to fully map out existing hardware profiles, and will be able to identify which systems can, and which cannot, take any update. In many cases, bringing a system up to capability can also be automated through the provision of other configuration changes (for example, new device drivers or firmware). An engineer being sent out to carry out remedial work on those systems that cannot be automatically updated will have full information of what is required – so minimising the need for repeat site visits where the engineer first has to find the root cause for the problem before ordering new parts, and reducing the cost burden of the human intervention.

  Full-featured automation systems not only provide a controlled means to roll out these changes, but also will have built-in rollback capabilities, enabling administrators to roll

back to a last known good release across all affected servers with a single push of a button. Indeed, many such tools can automate roll-back where the system can identify that a change has not been successful, raising an event in the process so that a systems manager can be informed that the specific device has not been upgraded, yet leaving the device in a working condition. Further, tracking such configuration changes in real-time ensures configuration integrity; an accurate representation of the infrastructure is always available, putting the user in control of configuration changes and reducing change-related downtime. Dynamic dependency mapping, depicting how components (applications, files and servers) interact, enables the full impact of any change to be made clear, and also leads to improved service delivery.

- **Patch Deployment**
  Many organisations have found that when it comes to server rationalisation, the presence of multiple instances of the same application can be a major problem. In many cases, the whole stack that an instance of an application is built on will have different patch levels than the same stack for a different instance. The only way to bring such instances together is to first rationalise the stack, and carry out sufficient retro-testing to ensure that these changes have not impacted the functional running of either instance. Only then can these instances be brought on to a single server or virtual environment. Data Centre Automation can speed this up through ensuring that both stacks are brought rapidly to the same patch level, and that these patches are not applied until the underlying servers are capable of accepting the patches through having the correct device driver versions, free resources and so on.

  Further, non-compliance with security patching can be costly, leaving systems open to attack from inside and outside the organisation. Enterprise level automation software provides operators with a real-time understanding of the state of all servers, their probable vulnerabilities, and likely patch sources. Such automation systems should facilitate the testing and validation of the proposed changes to each server, and then facilitate the installation and configuration of the patches on production servers. The systems should also allow users to assess the patch status across the organisation in order to demonstrate regulatory compliance, as well as have the ability to produce policy compliance reports.

- **Improve Business Service**
  By monitoring the status of business critical applications, and automatically moving them to another server in the event of a fault, an automated solution can dramatically increase functional availability. With advanced logic, not only can a “graceful failover” of the application stack be achieved, but failover policies set by IT can ensure that the best available server for a particular application is chosen at the time of failure, based both on application need and current state of resources.

Many organisations see DCA as purely a technical tool to help the IT department better manage the issues and problems that are perceived as plaguing such environments. However, DCA has a direct impact on how well an organisation can react to market forces – a well managed technical environment is, by its very nature, a highly flexible and responsive environment. A flexible technical environment is far better for meeting the changing needs f the business, and in being able to provide facilitation for changes in business processes forced on the organisation through internal or external forces.

6 What to look for in a DCA solution

As a result of its roots in basic provisioning and software distribution tools, just about all DCA software offerings should provide some basic functionality in this area, such as:

- Bare-metal server provisioning - the ability to remotely build new servers or re-provision old ones over the network without requiring direct physical access to the server. Typically this involves installing or reinstalling the operating system, application server (if needed) and required software components and applications and then, if possible, restoring data and settings.

- Configuration management - the ability to exercise detailed control over server hardware and firmware configurations as well as base application configurations. In addition to accurate and reliable automated configuration, the level of control provided by good configuration management provides the basis for smart asset management, patching, and automation itself.

- Root cause analysis – the capability to interrogate remote servers that are not performing as they should, and to identify what is causing the problem and so being able to address the issue should also be available.

- User provisioning – the capability to fully automate the addition and changing of user accounts, and to provide the needed security around such user accounts.

- DCA must provide automatic administration. By controlling almost all aspects of a server's configuration the need for administrators to directly manage servers is almost eliminated, removing the single greatest cause of outages - human error. Any administrative changes that are made can be audited and quickly backed out if necessary.

In the quest to provide integrated management of data center services, some vendors have spread automation of IT operations to the other parts of the infrastructure that are necessary to bring a server or application online. Other emerging trends, such as regulatory compliance, utility computing and ITIL/CoBIT, in addition to competition from traditional system management tools, have all spurred vendors to integrate additional capabilities.

- **Policy-based automation**
  While threshold-based automation is effective for simple web-based workloads - using triggers such as high CPU utilization to clone additional servers and reclaim them once the spike had passed - other applications are more complex to measure and scale up or down. Some DCA products have more sophisticated rules with graphical policy-building interfaces and the ability to deploy complex composite applications. Along with the assignation of application priorities based on business key performance indicators (KPIs), policies can be built that reflect what functions must be kept running – against those that it is only nice to keep running. This can form part of an automated disaster recovery solution - automated routines kick in upon a hardware failure and reassign remaining resources to critical applications.

- **Asset management**
  Keeping track of servers is an essential enterprise task that facilitates hardware and therefore software management, regulatory and license compliance, and security. DCA products have begun using discovery tools and produce detailed reports using their knowledge of hardware as well as the application packages that have been installed or removed. The level of granularity in many tools is high – for full DCA functionality, it is necessary to be able to recognise sub-assemblies (e.g. disk controllers/drives, network interface cards and so on), along with their associated driver details. Through this, DCA can advise as to whether patches, updates and full applications can be safely applied to the hardware resource, and on alternatives where necessary.

- **Discovery and interdependency mapping**
  Discovery and interdependency mapping tools enable DCA products to deal with unmanaged systems or unknown applications (i.e. systems not under their management, for which they do not have accurate configuration and audit data).

- **Audit and compliance**
  Using configuration data to detect drift from a desired state (a "gold standard" configuration, or a firm's IT policies); this capability may be part of the configuration management tool. Some vendors add connectors to service management systems

like HP Peregrine and BMC Remedy to differentiate between legitimate changes and errors, and so avoid bringing down a critical application by forcing a configuration back to a stored version.

- **Granular configuration discovery and consistency check**
  Some vendors have gone a step further than the above two points. Configuration drift, a common cause of downtime, can be caused in any of the three following ways:
  1) changes deployed to production environments are not also applied to the pre-production environment - the testing environment becomes invalid
  2) changes such as patches are not deployed consistently across clustered environments - leading to unreliable failovers
  3) configuration inconsistencies between production and disaster recovery environments can often render the DR environment ineffectual when needed.

  Administrators need deep visibility into such configuration inconsistencies.

  Granular Configuration discovery - provides not only discovery of servers, operating systems and applications, but also automatic discovery and tracking of specific detailed application parameters.

  Consistency Check capability provides the ability to detect and remediate consistency drift across different servers, applications or environments – from entire applications or file systems down to individual configuration parameters.

- **Workflow**
  In today's data centre operating environment, complex processes need to be executed consistently and efficiently to avoid costly re-work, delays and outages. A solid workflow capability enables both end user orchestration (process automation) and infrastructure automation (runbook automation) to streamlines critical data centre processes. It does so by coordinating notifications and providing user input between IT groups (such as DBA, Ops, and Sys Admin), enabling users to define and enforce IT standards while at the same time driving task automation both at the server layer and at the storage layer.

  The workflows of processes for dealing with particular situations (e.g. an outage, or the provisioning of a standard database server) can normally be encoded by means of a workflow modeller, using an easy-to-use graphical interface.

- **Network management**
  The ability to configure some aspects of the network - Ethernet switches, firewalls, and load balancers - is common. More sophisticated solutions include network device discovery and configuration capabilities.

- **Storage automation**
  Storage automation is a weak area for many DCA vendors due to a lack of standardisation in the area of storage management with respect to managing complex heterogeneous environments. Some vendors get round this by only catering for their own storage management tools – others do provide more heterogeneous offerings.

As well as the technical aspects of DCA, the following challenges can be addressed through the use of automation tools within an organisation, providing discrete value add and better flexibility in the manner the organisation can react to internal and external market forces:

- **Costly Operations**
  High support ratios (the number of administrators to maintain the number of servers) and over provisioned, highly decentralised IT environments can be rationalised, and the human resources can be better deployed in providing value add solutions to the organisation, rather than systems maintenance and fire fighting.

- **Poor Service Quality**
  The use of ad-hoc IT processes can lead to a high number of human errors, as well as constant change. Automation minimises or removes the ad-hoc aspect, replacing it with best-practice, and also removes the main areas of human error.
- **Security Concerns**
  The constant challenge of propagating and validating required security patches across large, heterogeneous environments is easily addressed through the centralising of patch management and the automation of the controlled roll out, remediation of underlying systems and implementation of the patches.
- **Time-Consuming Tasks**
  OS and application updates on hundreds or thousands of servers are time consuming, and require too many specialists. Automation removes the overall scale and resource issue, providing the capability to roll out updates to massive numbers of systems in a controlled manner.
- **Application Downtime**
  Unsuccessful and inconsistent changes to applications are a primary cause of system downtime and failures. Automation can ensure that any change can be handled by the receiving system before the change is implemented, that remediation is automatically carried out wherever possible, and that root cause is provided prior to any site visit where remediation cannot be automatically carried out.

At a business level, DCA not only offers technology savings, but also direct business value through faster response to the business' needs and a more flexible environment for the future.

7 Virtualisation, Dynamic Provisioning and Application Control

Virtualisation has opened up new vistas in DCA, in which servers can be treated as a pool of available resources that can be redeployed dynamically to meet changing workloads as dictated by the business' needs. This allows advanced workload management, based on demand for an application or function - if there's a spike in traffic, or performance degrades, the system is able to switch to another server, add another server to the environment, or to load balance to increase capacity immediately as required. Where a new server resource is required, the solution needs to be able to automatically configure all the necessary software components on a second server made available through virtualisation.

If automatic management of the underlying infrastructure was the original goal of DCA, then automatic management of the applications that run on it surely must be the ultimate goal. The combination of technologies offered by virtualisation and DCA makes it possible to manage applications in much the same way as servers. Once an application's base run-time requirements, such as its operating system and memory needs (e.g. runs on a specific version of Windows, requires 1GB of memory), network and storage connectivity, dependencies across internal application components and tiers, and its business priority are defined, then users can create and enforce policies based on those requirements to control when and where applications run across heterogeneous physical and virtual environments enabling them to maximise server utilisation, increase application availability and flexibly respond to changes in application workloads. IT operators can monitor, and start and stop complex applications across numerous physical or virtual servers, across every major operating system, in a secure, error-proof way.

From the above, it is apparent how complex this can be – and as we move to more dynamic environments, it becomes increasingly difficult to define the base needs and interdependencies and keep them all up to date manually. Automation is key – we need the tools to carry out all the initial discovery tasks, to maintain all the contextual information on the interdependencies within the system, and to provision, deprovision, patch and maintain the various layers.

8 Things to Ask a DCA vendor

Many DCA products are strong in one area, but weak in another. As a result, unless they are careful, customers may end-up having to buy multiple DCA products with overlapping capabilities in order to get all the capabilities they want. DCA product deployment costs can be high (it is worth noting that there will be no DCA tooling in place to help roll out the initial stages of the solution) so the usage of multiple solutions is something to be avoided if at all possible.

Here are some questions to ask yourself and any DCA vendor before investing in a DCA solution: -

- Is it comprehensive (in that it addresses all the core elements of data centre infrastructure - storage, physical and virtual servers, and networks) and is it an integrated solution, delivering capabilities in backup, storage management, server management, and application performance management across a heterogeneous data centre?
- Does the solution provide active management across all tiers in the infrastructure?
- Does the solution support all major application, database, server, and storage platforms?
- Does the solution support all major virtual machine platforms across all major operating systems?
- Does the solution discover all servers, applications, and any complex dependencies among them, and track any changes in real-time?
- Does the solution provide a granular level of visibility and control for virtualised environments by monitoring the applications within the virtual server, the virtual server itself, and the underlying hardware, as well as enabling the user to start, stop, and migrate the applications and the virtual servers across hosts?
- Does the solution enable control where multi-tiered applications run across heterogeneous physical and virtual environments?
- Does the solution have centralised, policy-based application management capabilities to help maximise the availability of critical applications?
- Does the solution provide patch and version management across the infrastructure?

The services that need to be provided by a DCA solution are complex by necessity, and therefore any solution needs to be comprehensive. Attempting to build a solution through utilising multiple best of breed solutions can lead to further complexity that needs to be managed. Quocirca recommends that as few a number of solutions as possible be utilised to meet the overall needs of the environment.

9 Conclusions

DCA is becoming increasingly important, as the complexity of the data centre and dependent systems external to the data centre itself grows. The advent of more utility-based computing architectures (SOA, SaaS, Grid) all point towards greater dynamics in the environment – and this will be further pushed by the business requiring greater flexibility and speed of response in facilitating the changes to its business processes driven by increased competition in the markets.

Any choice of DCA solution needs to be all-inclusive, as the move towards homogeneity within most organisations has stalled due to the increase in standards adoption making base level heterogeneity less of a problem.

DCA can also provide end user self service – the capability for users to carry out specific changes for themselves that may have had to go through a help desk function in the past. Besides the obvious cost savings that such self-service provides, the greater speed of response and end user buy-in helps to create a more responsive system.

Overall, Quocirca believes that a well-chosen DCA solution will not only cut costs within the data centre environment, but will free up IT resources (at both the hardware and human level) so that greater investment in value add projects can be made. Along side this, Quocirca also believes that DCA creates a far more flexible environemnt where the business regains overall control, being able to make process changes and strategic decisions knowing that the IT infrastructure will not be the constraint that it has often been in the past.

About Symantec

Symantec is focused on helping customers protect their infrastructures, their information, and their interactions. Headquartered in Cupertino, Calif., Symantec has operations in 40 countries.

Enterprises and consumers need to keep their infrastructures up and running 24x7. They need to be able to access information anytime and anywhere. That means that their critical systems must be up and running all the time.

Therefore, it is important that they protect the physical systems, the operating environments, and the applications – across all tiers of their infrastructure. They must protect a broad range of information types – from email to business documents to digital photos to audio and video files. And, they must ensure that the interactions – the connections, the collaborative environments, and the movement of data while in use – are protected.

Availability Technology Summary

Symantec offers products for backup and recovery of data and systems, optimizing storage resource utilization, simplifying administration of heterogeneous environments, and providing continuous availability of mission-critical applications and data.

Application Performance Management products optimize the performance and availability of enterprise applications.

Data Management Solutions are designed to protect, backup, archive, and restore data across a broad range of computing environments – from large corporate data centers to remote groups and desktop and laptop computers. The products integrate to provide solutions to manage data throughout its lifecycle – from creation to disposal, onsite and offsite, across all levels of storage hierarchy – including disk, tape, and optical storage media.

Infrastructure Management Solutions allow customers to manage virtually any function at any point in the lifecycle of their computing systems and devices. Solutions include network auto-discovery and IT asset management, operating system provisioning and application deployment, ongoing security updates and configuration management, rapid system recovery, de-provisioning, and help desk remote control.

Data Center Management Solutions help organizations gain control, improve service levels, and drive down costs within complex data center environments. These technologies deliver storage automation, virtualization, replication, high-availability clustering, server provisioning, and server management software across heterogeneous storage and server platforms.

For further information please visit www.symantec.com

About Quocirca

Quocirca is a company that carries out world-wide perceptional research and analysis covering the business impact of information technology and communications (ITC). Its analyst team is made up of real-world practitioners with first hand experience of ITC delivery who continuously research and track the industry in the following key areas:

- Business Process Evolution and Enablement
- Enterprise Applications and Integration
- Communications, Collaboration and Mobility
- Infrastructure and IT Systems Management
- Utility Computing and Delivery of IT as a Service
- IT Delivery Channels and Practices
- IT Investment Activity, Behaviour and Planning
- Public sector technology adoption and issues

Through researching perceptions, Quocirca uncovers the real hurdles to technology adoption – the personal and political aspects of a company'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 the adoption of the correct technologies at the correct time.

Quocirca has a pro-active primary research programme, regularly polling users, purchasers and resellers of ITC products and services on the issues of the day. 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, Dell, T-Mobile, Vodafone, EMC, Symantec and Cisco, along with other large and medium sized vendors, service providers and more specialist firms.

Sponsorship of specific studies by such organisations allows much of Quocirca's research to be placed into the public domain.

Quocirca's independent culture and the real-world experience of Quocirca's analysts, however, ensure that our research and analysis is always objective, accurate, actionable and challenging.

Quocirca reports are freely available to everyone and may be requested via www.quocirca.com.

Contact:
Quocirca Ltd
Mountbatten House
Fairacres
Windsor
Berkshire
SL4 4LE
United Kingdom
Tel +44 1753 754 838

quocirca