

## Ad-hoc Information and Intellectual Property The need for management and archival

### Contacts:

Clive Longbottom  
Quocirca Ltd  
Tel +44 118 948 3360  
[Clive.Longbottom@quocirca.com](mailto:Clive.Longbottom@quocirca.com)

Fiona Moon  
CommVault  
Tel: +44 118 951 6500  
[Fiona.Moon@commvault.com](mailto:Fiona.Moon@commvault.com)

*File servers continue to proliferate, and email has moved from a basic tool for exchanging documents to an inherent part of an organisation's communication and collaboration portfolio. With such ad-hoc information sources now being a mission critical function, the need for data management and archival is becoming critical.*

- **The growth of ad-hoc information within organisations continues to grow**  
Organisations are creating more and more information, the majority of which will never hit a formal storage system. Simply managing the growth of file server and email storage systems is taxing the majority of organisations. Effectively mining the value of such information is proving even harder.
- **The corporate importance of email and file storage is growing exponentially**  
Most organisations find that email is now being utilised as a form of information back up for many employees. Although the stateless capabilities of email servers means that this is technically an effective approach, the impact on storage volumes of having multiple copies of documents held on file servers and across multiple email storage systems is immense.
- **Both legal and organisational compliance needs are driving a requirement for better ad-hoc information management**  
Legal needs for process auditing, information tracking and disclosure are growing and changing constantly. Without the capability to keep track of ad-hoc information, an organisation's capabilities to demonstrate compliance will be effectively lost.
- **Organisational decision making is essentially impossible where the majority of information cannot be easily accessed**  
Making decisions against a sub-set of information is dangerous, yet most organisations can only effectively report against information held in formalised storage systems, such as databases and document management systems.
- **In the majority of cases, storage resources are woefully under utilised**  
Having dedicated storage systems for specific applications can lead to utilisation levels of less than 30% on storage systems.
- **Tooling is required that creates a virtualised storage resource pool**  
By choosing the right storage management environment, storage can be virtualised and logically partitioned to provide a highly flexible environment that is far more responsive to the organisation's needs and can provide more efficient means of utilising storage resources.
- **Flexibility and usability are key requirements**  
The tooling chosen must be easy to implement, must be inclusive in its support of existing and future storage hardware and must provide interfaces that are easy to use for both end users and technical staff alike.

#### REPORT NOTE:

This article has been written independently by Quocirca Ltd to provide an overview of the issues facing organisations in ensuring adequate management of their basic file and email environments. 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.

#### Conclusions

Ad-hoc information stores have historically been overlooked by organisations which have regarded simple back up and restore capabilities as being sufficient for business' needs. However, both internal and external pressures have been growing to make it that such stores have inherent value to the company, and that control must be exercised over the stores so that the full commercial value can be more easily uncovered and made available to enable better decision making.

An independent report by Quocirca Ltd.

[www.quocirca.com](http://www.quocirca.com)

Commissioned by CommVault

## 1. Introduction

Organisations started to store files centrally to get around the problem of unreliable client devices losing information on failure. By centralising ad-hoc (i.e. unstructured, as in reports, emails and other files) information storage in the same way as more formal (i.e. structured, as in data held within a database) data, better control could be applied to back up and restore, with the added benefits of better security and greater corporate visibility of information.

At the same time, email, which had started out as a basic means of exchanging information with others, rapidly grew to become a major form of collaboration and communication between people within an organisation. Today, email usage has spread outside the organisation to include suppliers and customers, and often comprises the exchange of a mix of formal and less formal information. Indeed, research by Quocirca has shown that email is one of the most utilised methods of carrying out transactions between businesses.

However, as the quantity of ad-hoc information has grown, the capability to effectively manage it has diminished, whilst the value of the intellectual property held in such files has exploded. That information now has to be stored for longer periods of time is compounding a problem that is just growing for the majority of organisations. Very few companies have a cohesive approach to dealing with ad-hoc data, splitting the management of files and email as if these two areas were completely separate.

Indeed, the overall business problem is often made worse as organisations try to address the problem – the capability for information exploitation is impaired, as the separate silos of information resources are further broken down into on-line, near-line and off-line silos, with few points of interaction between them.

This document looks at the need to regard ad-hoc data as a single corporate asset – and to manage it accordingly to maximise the overall value to the organisation itself.

## 2. Files and Emails

### Main Findings:

- Ad-hoc information should not be regarded as of low corporate value
- File servers and emails need to be regarded as a single resource – regarding each as separate entities will lead to problems with decision making down the line

The multiple informational silos caused by having multiple file servers and email data stores means that organisations will need to take an alternative approach to dealing with the problem for the future. As file servers get larger, performance slows, and back up and restore becomes more problematic. Domain searches, where information is searched across multiple servers, becomes too slow to be of use to end users, who simply stop using such functionality.

On top of this, as the usage of email has grown, it has brought other problems along with it, apart from just data volumes. Email provides an ideal way for outsiders to introduce viruses in to an organisation, and provides a means of accessing people in ways that would be difficult by other means, and so we have seen the growth of Spam. We have

also seen the growth of regulation on a world-wide basis, and this has meant that email has to be more controlled and has to be capable of being contextually positioned against other interactions with customers and suppliers, and within any corporate decision making process.

Quocirca has also found that many end users utilise email as an ad-hoc data back up system – sending copies of documents to themselves so that they are stored off their own machine, and are essentially accessible wherever the person is working, on whatever device.

A big issue that both file servers and email create is information duplication – multiple copies of the same file can end up being stored locally to the user, centrally via servers, and then on other servers and desktops as emails are forwarded. When combined with the inherent fragility of the majority of basic storage systems and of email systems’ “out of the box” storage approach, the dependence of many business processes and issues on uncontrolled file servers and email has to be addressed.

Further, we need to look at how the lifecycle of information changes – at what stage does data become information, at what stage does the usefulness of information decrease, and at what stage can we regard information as being of no value to the organisation? Within this, we also need to consider how the real value of the information is mined – how do we ensure that the results reported from a search covers all of the available information, not just the silos that were available to the search tool at that time. Organisations have to look to exactly how such information exploitation takes place through the provision of a single environment for both email and file storage.

## 3. Corporate, Vertical and Legal Compliance

### Main Findings:

- “Compliance” is not a set of different solutions – the aim should be to create a compliance oriented architecture, where information compliance is built in, not bolted on
- Compliance needs will be dynamic – and any chosen solution must provide the flexibility to allow for such dynamics

Uncontrolled file servers hold a lot of information that will never enter corporate higher end formal document management systems, such as Documentum or IBM FileNet. For organisations to depend on any level of corporate or legal compliance based on only monitoring, auditing and interrogating information held in formalised stores is a nonsense, and a dangerous one at that.

As well as the way that file servers have been regarded as base storage with little value, a major issue in email usage has been how email has been regarded as a separate system to the rest of an organisation’s intellectual assets. Although email can be called as a service from virtually any application via standardised calls, it rarely leaves any sign within the originating application. When we look at compliance, this can lead to major problems – a business

process may be owned and managed by an enterprise application, but all it needs is for a single employee to initiate an email for the process to be broken as far as compliance is concerned.

However, when we look at compliance within organisations, we need to determine an organisation's actual risk profile. Is compliance seen as being really necessary, or is it a risk that the organisation is willing to carry, as the financial and brand impact of being found to be non-compliant is outweighed by the financial cost of being compliant.

Compliance needs change over time – a large impact on compliance law is the political climate at any one time, and a change of government can lead to changes within the compliance laws. Without an underpinning of flexibility, any compliance “solution” implemented to deal prescriptively with a specific compliance need will inherently fail until it can be modified to suit – whereas an architecture where compliance is provided through dealing with granular information meta data tags, security and virtualisation of the underlying assets can be made to match the changes in needs as required.

File servers and email need to be managed as a peer environment to other more formal storage systems, becoming part of the overall virtualised pool of information that can be tagged and linked to be part of the over-reaching business process. Once we have this level of control, at the file, email and other storage level, the application of fine grain controls against interlinked intellectual property assets becomes possible, and we can begin to move to a Compliance Oriented Architecture, where matching information against process becomes far easier, and demonstrating compliance to specific needs becomes possible across the board. The main thrust for compliance should not be in being able to demonstrate that processes adhere to some rules written in the past, but in that all information sources and assets are fully audited and managed in a manner which allows any compliance need, whether this be legal (as in the EU Directive on Statutory Audit or MiFID), horizontal (e.g. FDA, CAA requirements), best practice (e.g. ISO9000, ISO27001) or process-based (e.g. ITIL, CoBIT)

#### 4. File servers and email as storage methods

##### Main Findings:

- The growing volumes of files and of emails, combined with the growth in value stored within them, means that a new approach to managing such assets is required

The use of file servers to provide centralised storage also means that standard back up and restore technologies are easier than when dealing with highly decentralised systems. However, as the usage of file servers has increased, we now have a situation where we have less centralisation than we would like, and a need to back up multiple different servers on a regular basis. With many organisation having poor policies in place, some file servers are not getting backed up – with the concomitant risks associated with the potential for losing data.

Also, the increasing usage of email as a stateless storage medium by users leads to various problems. Firstly, as email stores tend not to be based on single stores, with copies of

each email proliferating across multiple servers and clients, we have to look at backing up redundant information on a regular basis. Secondly, the basic storage mechanisms of email systems tend to be monolithic, making it difficult to manipulate individual emails from anything but the email client. Thirdly, built-in email archival tools are based on an off-load capability, focusing on creating a smaller primary database with few capabilities to manage the archived data at a granular level.

By treating email as a standard information feed, it is possible to create a far more efficient and effective system, where emails are stored once and once only within any single domain, and where information archival is carried out in a far more manageable manner.

Again, we have to look at file servers and email as being a common issue – we are trying to manage the corporate intellectual assets in a manner that protects the organisation. We need to ensure that the exploitation of the available commercial value within the information is more easily uncovered by anyone within the organisation. Only a fully rounded approach can make this happen.

### 5. Information Archival and Management

##### Main Findings:

- The effective tiering of storage resources creates immediate value to an organisation
- Storage virtualisation creates a single resource pool that provides technical flexibility and better supports business needs

Whereas file servers create unmanaged collections of files written directly to a magnetic disk, email systems create masses of data that tends to be held in multiple different types of data stores in multiple different locations. Any tool that enables better control of the file and data stores, creating a more formalised method of storage and retrieval, will provide both process and monetary value to a company.

Hierarchical systems are needed, where the information that has the most usage is stored on “tier 1” storage (expensive, high speed disk), being aged through to lower speed disk and then on to near-line and off-line storage devices as the need to access the information becomes less. This has to be linked with the legal and other needs of the information, such that compliance can be easily demonstrated when needed. A further complication arises as we need to look at the lifecycle value of the data and information – in general, information has more value towards the beginning of its life, tailing off as time goes on. However, this may not be the case in areas such as research, where initial findings are of little value until everything has been checked and is ready for patenting/publication, or in marketing campaigns, where iterations of information coming back from campaigns can have various levels of value throughout the campaign itself. To maximise the optimal utilisation of the underlying physical storage assets, a full understanding of such matters needs to be gained, and then codified into flexible business policies.

To create a flexible storage underpinning, the underlying storage assets have to be made available as a single resource pool, through the use of storage virtualisation.

Storage and virtualisation come together to provide a solution that is generally far greater than the sum of the parts. Whereas the basic concept of virtualisation is to make multiple different resources appear to be the same, one of the strengths of storage virtualisation is to make such a pool appear to be what it is not.

For example, we may find that backing up our main file and email stores is becoming too time hungry, and that we are continually backing up data stores to physical tape. However, we could take the approach of assigning a portion of high speed disk to appear as if it is tape, so enabling a very rapid back up to take place, and then we can “trickle” this from the disk to the tape as the storage and network bandwidth allows.

Virtualisation of storage also enables systems to be far more easily migrated and managed – for example, by abstracting the email system from the physical storage, we can move from one version of the email engine to another without the need to migrate the stored information from one physical media to another, yet we can also increase allocated storage without the need for the application to be taken off line at all.

## 6. High level approach

### Main Findings:

- A single storage engine approach provides greater opportunities for information exploitation.
- Business continuity and information security can be massively improved through a common storage approach.

When looking at putting in place a suitable file and email management and archiving solution, we begin to see that there is a strong need for a more holistic approach. File and email storage cannot be seen in isolation from other information storage, and we have to ensure that all non-formal intellectual property assets can be dealt with in the same manner as other, data assets, such as formal database stores and enterprise document management systems.

We therefore need a single underlying engine that manages all of our storage needs, with different aspects managing different types of storage.

Within this engine, it is important to ensure that the following capabilities are present:

- **Policy engine**

The capability to deal with the information according to the needs of the organisation and the law, either through looking at the security needs of the information, the storage lifetimes, the access needs of the information and so on.

- **Content granularity**

Only by providing a good level of information granularity can the correct policies be made to work. Each item of information must be dealt with as if it was a completely separate intellectual property asset, so that the correct security and lifecycle management policies can be applied against it.

- **Virtual and real HSM**

Just having the capability to cascade information from one physical type of storage to another is no longer good enough – we have to be able to cascade information through the system within the time and physical needs of the organisation. Therefore, there is a strong need to enable “logical” storage units to be set up – for example, high speed disk being made to look like tape – so that areas such as back up can be carried out far more rapidly.

- **Snapshot back up and restore**

For business continuity, the immediate need is for a restore system to be as up to date as possible. The capability to take rapid snap shots of data means that organisations can ensure that data back up can be far more timely than existing tape-based back up systems, and through the use of snap shot restoration, can ensure that the business is back up and running again as rapidly as possible after any data loss.

- **Ease of use – for users as well as techies**

Information management should no longer be an area where only technical people can provide solutions. There is a strong need to enable end users to have a self-service approach, recovering files and emails from their own back up environment as necessary, so cutting down on the cost of expensive technical support, and improving their own productivity.

- **Security of content**

Information security has to be built in, rather than bolted on. By utilising optimum levels of information granularity combined with information tagging, information security can be more easily managed in accordance with the organisation’s policies and legal needs.

- **Information Exploitation**

For the majority of companies, the only information that is effectively acted upon is that which is under formal control, for example formal data within databases, or ad-hoc information held within document management systems. Through the usage of a single underlying storage engine, users can exploit more of the available information directly – and so can be far more effective and efficient within their jobs.

## 7. Conclusions

**Main Findings:**

- Organisations need to look at how best to realise the commercial value of the information assets they have
- Email and ad-hoc files need to have a common approach in order for the overall corporate information base to be exploited to its full extent

In conclusion, organisations are in a position where the underlying value of ad-hoc information is not being fully realised due to the poor control of the storage of such assets,

To rectify this situation, it is necessary to review how ad-hoc information is being stored, and to create an inclusive environment where emails and files can be regarded as a single resource. Through this approach, information searches become far more effective, and the commercial value held within the basic electronic information stored by an organisation can be exploited to its greatest extent.

Storage approaches where a single engine underpins the available storage resources can create a highly effective virtualised environment which not only provides the means for more effective information management, but also the capability for employees to more effectively make decisions based on identifying and working with all information that is important to their task in hand.

Overall, Quocirca believes that the proliferation of uncontrolled information within organisations is now creating a major block to the overall commercial effectiveness of all concerned. As compliance needs continue to grow, and as information reporting becomes more important, the problem must be addressed. The ultimate aim has to be to create a virtualised, single resource of available information that is secure, allows for business continuity, that makes the greatest usage of the available storage resources and is simple to implement, manage and use. Flexibility will remain a major requirement, and any chosen solution must be able to embrace new storage systems and technologies into the existing environment, and provide suitable levels of management granularity for individual information assets to be effectively utilised and their inherent value to be maximised.

## About CommVault

A singular vision – a belief in a better way to address current and future data management needs – guides CommVault in the development of Unified Data Management® solutions for high-performance data protection, universal availability and simplified management of data on complex storage networks.

CommVault's exclusive single-platform architecture gives companies unprecedented control over data growth, costs and risk.

CommVault's software was designed to work together seamlessly from the ground up, sharing a single code and common function set, to deliver superlative Data Protection, Archive, Replication, and Resource Management.

More companies every day join the thousands who have discovered the unparalleled efficiency, performance, reliability, and control only CommVault can offer.

Information about CommVault is available at [www.commvault.com](http://www.commvault.com).

CommVault's corporate headquarters is located in Oceanport, New Jersey in the United States. The EMEA Regional Headquarters is located in Reading in the UK, and the Oceania office is based in Sydney Australia.



## 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 in the following key areas:

- Business process evolution and enablement
- Enterprise solutions and integration
- Business intelligence and reporting
- Communications, collaboration and mobility
- Infrastructure and IT systems management
- Systems security and end-point 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
- Integrated print management

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, 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 at no cost. Quocirca's reach is great – through a network of media partners, Quocirca publishes its research to a possible audience measured in the millions.

Quocirca's independent culture and the real-world experience of Quocirca's analysts 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](http://www.quocirca.com).

**Contact:**

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

## Managing Intellectual Assets A new approach to Compliance in Finance Markets

### Contacts:

Elaine Axby  
Quocirca Ltd  
Tel +44 20 8874 7442  
[Elaine.Axby@quocirca.com](mailto:Elaine.Axby@quocirca.com)

Clive Longbottom  
Quocirca Ltd  
Tel+44 118 948 3360  
[Clive.Longbottom@quocirca.com](mailto:Clive.Longbottom@quocirca.com)

Fiona Moon  
CommVault  
Tel: +44 118 951 6500  
[Fmoon@commvault.com](mailto:Fmoon@commvault.com)

*The increasing impact of local, regional and global compliance rules on the financial markets cannot be overstated. The need to identify, control and manage formal and ad-hoc information assets is growing – and has to be built on a suitable storage infrastructure.*

- **The compliance burden is growing and shows no signs of abating**  
Significant requirements have been imposed in recent years – SOx, Basel II etc – and new EU requirements (e.g. MiFID) this year further update the financial services regulatory framework and encompass more and more products and organisations.
- **To date, approaches to compliance have treated each requirement as a separate entity**  
This has meant many different ‘point’ solutions for compliance, and therefore a need to check each against the others to ensure that one view has not distorted another and that there is one version of the ‘truth’. As the regulatory burden changes, further individual solutions will exacerbate this problem, adding to cost and complexity.
- **What is needed is to treat information as a corporate asset**  
A single piece of underlying data should only exist once, but should be capable of being used in different applications – this would bring significant benefits from an overall business as well as a compliance perspective
- **A ‘Compliance Oriented Architecture’ presents a way forward**  
Such an architecture should be able to identify and capture information at the point of creation, should be able to apply security and tags based against a corporate policy and must be able to index the information to make identification easy.
- **Such an approach would bring significant business and IT benefits**  
It would enable compliance to be demonstrated more quickly and easily as well as enabling the business to better respond to market pressures. Treating information as a “service” means that performance can be measured in terms of the service levels delivered, increasing the focus on a service-based culture within the organisation.

### REPORT NOTE:

This article has been written independently by Quocirca Ltd to provide an overview of the issues facing organisations in the financial markets around the management of intellectual property assets. 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.

### Conclusions

Compliance can often be seen as an intolerable burden; building another compliance application and ensuring its consistency with existing solutions is time-consuming and expensive. However, with building the appropriate architecture, data can be captured only once, tagged and indexed and used by a range of applications. Rather than seeing data management for compliance as burden on the organisation, changing the organisation’s mindset can present an opportunity to improve not only compliance but wider business performance.

## 1. Introduction

We live in an ever more regulated world – and the financial sector has some of the most stringent information compliance laws around. From the same data protection laws that apply to all organisations, through to local requirements, such as those placed on the financial sector in the UK by the Financial Services Authority (FSA), and regional requirements such as the EU's upcoming Markets in Financial Instruments Directive (MiFID), the need to have full control of informational assets has never been stronger. From a process perspective, there is an increasing need to demonstrate that best practices are being followed – standards requirements such as ISO 27001 and 17799 (both related to information security management), Six Sigma and ITIL/CoBIT (related to general business and IT good practice), mean that a comprehensive and coherent data management system must be in place from both a business and IT perspective.

Financial services organisations have seen major change in recent years: market liberalisation has led to an explosion of new product availability, companies moving into new markets and extensive merger and acquisition activity. The impact on the IT systems of the organisations involved has been immense, and the response has typically been to build separate applications, each on its own platform, to support business expansion and meet the increasing legal and regulatory burdens.

Underlying these systems however are the core customer data: who are they? what are they buying? when did we last communicate with them? what did we agree to? what did they agree to? The problem comes when trying to access this data for different purposes, and come up with a coherent view of the data. Information has tended to sit in 'silos', with a different application being built for every compliance need that comes along.

In this paper, however, we try to illustrate that it is possible, by taking an integrated approach driven via corporate policies, to construct a solid information asset management platform – and create the flexibility to manage Compliance and other information requirements as they evolve.

## 2. The challenges of managing information assets for Compliance

### Main Findings:

- Compliance take many forms and the burden shows no sign of abating
- The risk of different 'point' solutions is that full compliance becomes difficult to prove, each solution working against its own data set.

Within the finance sector, the need to manage information assets to meet compliance requirements is becoming ever more acute. Some of the key compliance requirements include:

- Data Protection: data must only be used for permitted purposes, and must only be kept for a "reasonable" amount of time

- Sarbanes-Oxley (SOx): Even at a basic level, all EU based companies that deal with US companies at any level will need to be able to demonstrate SOx 404 compliance, while many will have to demonstrate greater adherence to the full SOx legislation
- regular and ad-hoc financial authority reporting: the need to demonstrate adherence to financial legislation, including fair treatment of customers, report on transactions and risk exposure and to report on changing market forces requirements
- BASEL II: implemented in the European Union through the Capital Requirements Directive (CRD), this sets out minimum capital requirements to meet various types of risk and requires firms to publish details of their risks and how they manage capital and risk, and financial institutions must be able to show all supporting data as required to demonstrate that the institution meets the minimum requirements.
- MiFID: an upgrade of the existing Investment Services Directive, MiFID updates and expands the European financial services regulatory framework. Different types of product and organisation will be included in the new framework. Due to be fully implemented in November 2007, compliance is already being worked on by the majority of finance institutions.
- Standards/best practice requirements: the need to show that "best practice" processes are being utilised along value chains and within a financial organisation. Although not a necessity, such compliance helps to breed a degree of trust in the customer and with other financial institutions that may be part of the overall solution. The more institutions that take up this sort of compliance, the more important it becomes to those yet to take the plunge – being seen as an untrustworthy financial source is patently not a viable approach.

These requirements are undoubtedly going to change and be added to in future years: SOx was a political response to certain issues seen in the US, and has already changed markedly, MiFID is a response to changes in financial markets and will almost certainly need to be changed further to reflect future changes in the markets, and statutory requirements such as Freedom of Information, currently imposed on public bodies, may be expanded to take in commercial organisations.

Alongside all of this is the need to manage the needs of the business and to be able to report against these – internal compliance, if you like. This will differ from company to company, yet has just as much validity when it comes to the need for a compliance solution as the legal requirements.

Many organisations have taken a piecemeal approach to these requirements, with solutions being implemented for the various Data Protection Acts applicable for multi-national organisations, for demonstrating the various levels of SOx compliance that may be required, for ISO standards compliance against processes and so on. The problem is then shifted on to the demonstration of complete compliance – each solution will have been built to work against specific data sets and processes, often creating its own data sets to do this. Further, with each solution working against its own rules and view of the "truth", full disclosure can mean having to run reports against multiple solutions and then correlate

and aggregate the results before making them available to the powers that be.

Running a report to satisfy a specific compliance requirement where the reach of the report is not fully controlled raises a possible further issue – not only does this run the risk of providing insufficient information to meet the regulatory need, but it could also provide too much information, in which case it may result in being out of compliance with a completely different need. As an example, the need to demonstrate compliance against, say, an ISO process may involve a high degree of personal information being accessed and dealt with. The ISO process has no interest in the data itself, just that the overall process has been followed. Now, if the data accessed is made visible to the person running the ISO audit (which has no impact on the legal standing of the institution itself), then it can be taken as read that it is likely that the data made visible will have broken the local data protection laws. Therefore, something which started as a means of demonstrating trust, not only ends up breaking that trust model, but also in moving the institution into being legally out of regulatory compliance.

Although the majority of institutions would find themselves in this position only by accident, the stance taken by the statutory powers is increasingly that this is no defence. It is therefore important that an institution not only understands all the compliance issues it is up against, but that it also creates strong policies that cover all the various aspects – and puts in place the correct tools to make the policy work.

### 3. The need for an information policy

**Main Findings:**

- Underlying all compliance is the same data, which needs to be appropriately treated so that it can be used in different compliance situations
- This data can also be used in other ways, for example, within sales and marketing

A financial organisation needs to have a strong information base to work against. What needs to happen is that single items of data are appropriately indexed and categorised so that they can be easily identified and contextually utilised in different compliance and other information situations.

Take, for example, a mortgage sale. The data related to this sale:

- will be needed for regular regulatory reporting to the institutions such as the FSA
- will need to comply with data protection legislation
- will form part of the company’s internal reporting, capital requirements and risk analysis – and be used for regulatory reporting on these issues
- will need to comply with consumer contract legislation
- will need to comply with legislation related to mis-selling of financial products.

- will need to be stored for a period of time exceeding the lifetime of the contract itself – maybe 75+ years
- will form part of a “value chain” of information with solicitors, estate agents and intermediaries and may need to be passed to other organisations such as external financial checking systems and possible court judgements, such as County Court Judgments (CCJs) in the UK
- may need to be presented in an anonymous manner to demonstrate different forms of compliance to other needs (e.g. ISO)

It is easy to extrapolate this problem to other financial areas, such as the trading of stocks, shares and other financial products, insurance services, loans and credit cards, and the managing of a whole range of products and investments for companies or individuals.

As well as the compliance issues mentioned above, there is an increasing need to pull together customer accounts and records (e.g. current/savings accounts, insurances, other accounts held within the same family) so as to be as responsive to their needs as possible – within the regulations. The organisation must be able to respond quickly to market dynamics – whether this is the launch of a new credit card offer, a change in interest rates, a new means of wealth management being launched by the competition and so on..

In the end, organisations need to take a different approach to data management. Below, Quocirca suggests a ‘Compliance Oriented Architecture’ (COA), which enables a horizontal view of the data required for compliance, enabling organisations to make all of its data available for compliance applications to be built as needed.

### 4. Creating a Compliance Oriented Architecture (COA)

**Main Findings:**

- Organisations should be seeking to implement an IT infrastructure that enables compliance.
- Information can be provided as a “service” to the business, indexed and tagged so that it can be used in a variety of situations

A compliance oriented architecture aims to make compliance part of the infrastructure, rather than a bolt on solution. To do this, it is necessary to move away from looking at each type of compliance as a separate entity, and to look at how the individual informational assets have to be addressed.

For example, within the mortgage example outlined above, the various items that make up a full view of the mortgage offer itself will fall under different compliance requirements at different times.

Taking a vertical view of compliance, by introducing DPA solutions, followed by ITIL solutions, followed by MiFID solutions and so on, runs the risk of creating a perception of compliance, where when it comes to the crunch, we find that each successive layer of a solution has, in fact, broken one or more of the preceding layers.

Also, with much of the compliance burden being politically driven, we can expect the rules underpinning such compliance to change as the political wind changes. If, however, we look at how we can identify and tag information at the point of creation in such a way so that we can interrogate the whole environment in an open yet secure manner, we find that we create a highly flexible and responsive environment where compliance is built in, rather than built on.

This does require a cohesive and coherent approach to how we look at the data – data has to be regarded as a corporate asset and must be virtualised in a manner that makes the information available as a service to any item that requires it. To ensure that the information can only be utilised by the calling services that are allowed to access it, security has to be applied at a highly granular level. To make the identification of the information easier, information must be easily tagged and indexed.

This dictates that the underpinnings be based around an engine that is there as a service itself, an engine that comprehensively deals with information on an end-to-end basis. This engine should be able to identify and capture information at the point of creation, should be able to apply security and tags based against a corporate policy, must be able to index the information to make identification easy, and must be highly responsive to the needs of the business itself.

Through the use of such an informational engine, we can then start to build our compliance oriented architecture. As the underlying information is now available as a service, it is possible to create composite applications that can be utilised where we need to demonstrate compliance against a specific requirement.

For example, should there be a need to demonstrate DPA compliance for a specific country, a composite application can be created that rapidly identifies the information that is required to demonstrate compliance – and only that information. Information that falls outside the compliance need is not presented to either the external compliance or

internal compliance officers. Should there be a requirement to demonstrate adherence to ITIL processes, a similar composite application can be created that shows the processes and the information assets that have been utilised here.

The key is to ensure that only the information assets that have to be disclosed are disclosed – and this can only be done by ensuring that all the information is callable, but that the granularity of meta data and of security ensures that only the “right” data is made visible.

Through the use of a COA, not only can compliance be demonstrated, but internal reporting is massively improved, leading to faster and better decisions, based on more of the available information. A COA provides information as a “service” and therefore the organisation has greater potential to measure and benchmark its Service Level Performance as judged by the business requirement – how quickly did we get data? was it the right data? – rather than the storage based IT metrics. These added advantages of a COA mean that whereas a pure compliance approach is often seen by the business as a cost to the bottom line, the implementation of a COA can be viewed as an information investment, with a distinct value to the business.

## 5. Conclusions

Compliance can often be seen as an intolerable burden; building another compliance application and ensuring its consistency with existing solutions is time-consuming and expensive. The compliance burden shows no sign of abating, so a new approach is needed.

By building a “Compliance Oriented Architecture”, data can be captured only once, tagged and indexed and used in a range of applications. Rather than seeing data management for compliance as burden on the organisation, changing the organisation’s mindset can present an opportunity to improve not only compliance but wider business performance.

## About CommVault

A singular vision – a belief in a better way to address current and future data management needs – guides CommVault in the development of Unified Data Management® solutions for high-performance data protection, universal availability and simplified management of data on complex storage networks.

CommVault's exclusive single-platform architecture gives companies unprecedented control over data growth, costs and risk.

CommVault's software was designed to work together seamlessly from the ground up, sharing a single code and common function set, to deliver superlative Data Protection, Archive, Replication, and Resource Management.

More companies every day join the thousands who have discovered the unparalleled efficiency, performance, reliability, and control only CommVault can offer.

Information about CommVault is available at [www.commvault.com](http://www.commvault.com).

CommVault's corporate headquarters is located in Oceanport, New Jersey in the United States. The EMEA Regional Headquarters is located in Reading in the UK, and the Oceania office is based in Sydney Australia.



## 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 in the following key areas:

- Business process evolution and enablement
- Enterprise solutions and integration
- Business intelligence and reporting
- Communications, collaboration and mobility
- Infrastructure and IT systems management
- Systems security and end-point 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
- Integrated print management

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, 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 at no cost. Quocirca's reach is great – through a network of media partners, Quocirca publishes its research to a possible audience measured in the millions.

Quocirca's independent culture and the real-world experience of Quocirca's analysts 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](http://www.quocirca.com).

**Contact:**

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

The logo for Quocirca, featuring the word "quocirca" in a lowercase, sans-serif font. The letters "qu" are in blue, "o" is in red, "c" is in blue, "i" is in red, "r" is in blue, and "ca" is in red.

## Data Management Today

### An overview of data management in the light of changes in enterprise storage infrastructure

#### Contacts:

Dennis Szubert  
Quocirca Ltd  
Tel +44 845 397 3660  
[Dennis.Szubert@quocirca.com](mailto:Dennis.Szubert@quocirca.com)

Fiona Moon  
CommVault  
Tel: +44 118 951 6500  
[Fiona.Moon@commvault.com](mailto:Fiona.Moon@commvault.com)

*The storage landscape has become increasingly complicated in recent years. Along with the rapid adoption of network storage technologies and tiered storage models has come a multiplication of replicated data copies. How well have storage management tools coped with the changes?*

- **Organic growth of storage systems has led to massive disconnects in many organisations**  
Silos of data and of data management tools have grown as multiple solutions within organisations have grown. However, as organisations have attempted to rationalise and integrate the applications, the data has often been left in a chaotic state.
- **Many organisations view all data as having the same value**  
A lack of understanding between the concepts of “data” and “information” means that much data of low value is being stored multiple times on expensive enterprise storage systems.
- **Data in itself has little value to the business**  
Data needs to be turned into information – from information comes knowledge, knowledge helps decisions to be made. A data and information management strategy needs to enable this process.
- **Much data is held insecurely against the needs of the business**  
The lack of information policies, and the tools to enforce such policies, is creating security loopholes that many organisations are unaware of.
- **It is time for a review**  
The management of data as information assets requires a review of an organisation’s approach to storage. Tinkering with existing systems may not work, and may make matters worse over time.
- **An integrated approach is required**  
Data and information systems that are tightly integrated, based on a single code base will offer the best opportunities to create a solution that meets the business’s needs.
- **The key areas for data and information management must be supported**  
Data discovery, classification, policy enforcement, security, data backup and restore, archiving, mirroring, audit and secure deletion (purge) should all be provided.
- **Abstraction allows for greater flexibility**  
Moving the data away from being fixed to linkages with the underlying physical assets provides a basis for greater flexibility, enabling data migrations and hardware upgrades that cannot be carried out otherwise.

#### Conclusions

As volumes of data continue to grow at a rapid rate, new approaches to data and information management are required. Ensuring that chosen systems provide a high level of function while allowing ongoing flexibility and durability is key.

#### REPORT NOTE:

This article has been written independently by Quocirca Ltd to provide an overview of the issues facing organisations in the financial markets around the management of intellectual property assets. 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.

## 1. Introduction

*Over the past few years, enterprise storage infrastructures have undergone a major transformation, bringing with it a whole new set of operational challenges. The traditional notions and requirements for data management are no longer valid, and a different approach is needed to address these new management challenges.*


*Data management has not kept up with the changes that have occurred in storage technology. Customers today need an integrated solution for administering heterogeneous data storage systems that makes their hard-pressed administrators more productive. In short, data management needs fixing.*

## 2. Data/Information Management

### Main Findings:

- There is a need to understand the difference between data and information management
- Without a top-down DM / IM strategy it is difficult to manage and maintain information

Figure 1 shows an information triangle – an organisation can pull out information from its underlying data assets, either through running reports against structured data stored in databases in the creation of textual reports, or on the aggregation of views. From this information, decisions can be made with a degree of knowledge that did not exist previously.



**Figure 1**

The value of data is not always fully understood within organisations and perhaps because of this, classification of data is not widely adopted (except perhaps where it is defined by a physical storage model i.e. what devices the data is held on) or data writer type (e.g. a customer database). This holds true not only for on-line but also near line and off line data sets. This is an important point as a proportion of current and the majority of all historical data can be found on storage media that resides away from the day to day running of the business on primary storage devices. It is also important to recognise that without having unified access to all of an organisation's data sets then the ability to extract firstly information and secondly knowledge from the infrastructure becomes a lengthy and costly process.

It is well understood that most storage subsystems are underutilised, yet few organisations are able to accurately detail existing storage utilisation (both online and offline) and are therefore unable to trend, based on current and predicted growth rates.

Alongside this, information and / or data tracking are not widely adopted processes, leaving many organisations with little capability to measure the cost of operations in supporting a data management / information management infrastructure.

For many organisations, a fragmented deployment of data management tools has resulted in an inability to extract valuable information from its infrastructure and, worse still, the inability to provide simple business services such as: data discovery, audit, search, retrieve and purge.

## 3. Identifying the Problem

### Main Findings:

- Storage has been segregated, and the storage management systems have followed suit
- Organisations should stop treating their data assets as simple data and start considering them as information, or intellectual property
- Not all systems, information and data are of the same value and as such organisations should look to ways of classifying them.

In approaching this problem, it may be helpful to ask some simple questions:

- Why is the storage underutilised?
- Why are there so many duplicate copies of data? Where are they and what are they used for?
- Is it possible to predict accurately storage requirements over a 1 – 5 year period?
- Is data easily classified? If not why not? And at what level should classification be carried out?
- Should all data be managed at the same level?
- What information can be derived from the existing data sources?

Data has become segregated both by IT divisions and organisational structure, and the management systems as well as the types of storage have likewise been segregated. It is not uncommon to find multiple storage management tools in different areas of operation within one organisation, leading to duplication of people (skills and effort) and wasted resource in terms of both licensing and management.

Organisations need to stop treating their data simply as that: blocks of data, but should start considering it as information. Information has a business value, owner(s), a location, a need for information retention that needs to be driven by policy, and should be tracked, accessed and reported on regardless of location on physical or logical storage subsystem (or even its geographic location).

If data issues are not dealt with, competitiveness in the market will be hit, as those fleetier of foot and utilising more

flexible information management tools will be able to respond to changes in the markets faster. By taking control of data and information issues as soon as possible, the organisation can be at the forefront of the market, and can more rapidly respond to the demands being placed upon it by its customers and competitors.

## 4. Underlying Business Issues

### Main Findings:

- Maximum benefits are obtained when solutions are employed in concert
- Integrated solutions ensure seamless operation

Specifically relating to data and information, IT is often viewed as being unable to comply with the business expectations for:

- Maintaining and controlling information growth.
- Data and information recovery (both following a loss and also for audit or analytical purposes.)
- Operational cost capping.

On the last of these, a widely quoted assumption is that the cost of management to the cost of the storage is in the ratio of 4:1. If that is correct, then businesses are facing an ever increasing and potentially non-linear compound problem, as storage requirements continue to grow in an uncontrolled manner.

As a result of this perceived failure to respond to the business, IT is often considered as a “tax” on the business - it is often seen to offer little more than an insurance policy in the event of a problem, and often an insurance policy that doesn’t pay out when problems do arise.

As organisations strive to move away from the cost centre model widely deployed today to service orientated architectures, it is imperative that the direction and strategy for data and information management be managed from the highest level and information stored in the form of data should be recognised as the corporate asset that it is.

That data is not normally viewed in this way can be demonstrated by asking staff the value of what they are storing. They may be able to identify the item in question as a database, a document or a mail file, but the vast majority of operations staff have no concept of the value of the information in itself. As a result, there is often no concept of differing support models to support these information types. Yet, it seems reasonable that not all data (e.g. backup, archived, replicated etc.) should be treated in the same way, as different types of data will have different levels of value to the business. For example, information held regarding a patent application will have far more value than information being held against a health and safety requirement – but both will need managing.

This lack of understanding of the value of information throughout IT departments (and many parts of the business) hinders the ability to make sound financial decisions regarding online, backup-to-disk, retention, archive and snapshot storage both in trend analysis, and quantitative purchase.

Another (and potentially larger) problem looming regards data retention for legal or audit purposes. Without knowledge of the underlying data value, defining some standard retrieval scenarios would quickly expose the inability for organisations to retrieve information as potentially required by law, or to report effectively within the organisation as required for corporate governance.

For many organisations, continuing with current storage management strategies is not viable – they will fast become ineffectual in providing responses to requests for information in context and in the process of trying, will lose competitive edge as all resources including much of the financial resource will be mobilised to manually deal with this new requirement.

It’s time to stop, take stock of the situation (with both current and future expectation in mind) and put some significant effort into developing an information strategy that meets the business requirements on an ongoing basis.

## 5. IT Issues

### Main Findings:

- It can be difficult to define a top-down DM/IM strategy
- Data segregation negatively impacts storage purchasing and ownership

In the majority of cases, it is seen that there is an inherent inability to manage and control data consumption across the enterprise. This is due to a number of factors:

### Difficulty in defining a Data or Information Management Strategy.

This difficulty stems primarily from an ownership model. In defining an IM strategy, a holistic view of all an organisation’s data must be taken, such that ownership, value, retention and other classification objects can be assigned “top down” to the separate business and IT functions. In the absence of this encompassing approach, each IT function has strived to do the best with the tools, people and vendors to whom it has access.

### Segregated IT Divisions.

Typically the organisation’s data is subdivided (as is the ownership) across several systems and line of business units, with different applications, operating systems and databases causing issues at the bottom up, and different business units, different processes and different business requirements causing issues from the top down.

### Purchasing.

In many organisations, a lack of coordination between functional groups can lead to one or more of the following procurement issues:

- Lack of standard buying programmes aligned to match IT strategy.
- Lack of standard purchase policies (e.g. multiple backup tools)
- One rule for one division, and another for another division.

It is imperative that purchasing departments are involved with the DM / IM strategy, to ensure that such issues can be dealt with up front, and the maximum benefits received from a standardised toolset strategy.

#### “The lions’ share spend factor”.

When making decisions regarding which storage management portfolio to review, it often falls to the vendor who has sold the most storage (generally measured in financial value) to the organisation. On the one hand, this has practical implications such as those around support models and relationship strength. On the other hand, this is a potential pitfall for which there may be no remedy:

- It becomes difficult to become hardware agnostic.
- It is nearly impossible to use vendor leverage to obtain best pricing.
- All of the eggs are in one basket, with an inability to flex as business requirements on IT change.
- Decisions are being made on a price / flexibility basis and not necessarily one that meets an organisation’s data management strategy.

#### Data Segregation.

Customers have been found that data segregation has happened because there’s been no better way. But the result has been:

- Costly, wasteful storage rationalisation programs (often driven by hardware vendors)
- A requirement for bespoke data migrations
- A lack of cross-business data ownership – instead data is pooled by OS type (Microsoft, Unix, Mainframe), application type (Exchange, Oracle etc) or storage type (SAN, NAS, or even by vendor)

Through being aware of these pitfalls, an organisation should be better able to prepare itself for a more strategic approach to data and information management.

## 6. Technical approach

#### Main Findings:

- There is a strong need to abstract the storage and data pools from the underlying physical assets
- It is recommended that a common code base be utilised to ensure commonality of management and control

Organisations have attempted to contain data growth while maintaining a short leash on management costs and capital outlay, yet without much success in many cases. Primarily, this has been because of complex vendor integration policies, combined with a lack of availability of highly functional, vendor agnostic tools. There has been a discrepancy between code and vision from storage management vendors that have expanded their portfolio primarily through product acquisition; these products have not always been designed to interoperate from conception to market deployment. Also, many management systems have been unable to cope with

market shifts with technology (such as virtual tape libraries (VTL), and backup to disk (B2D), etc).

As with many areas of IT, simplicity is the key.

If two simple rules are followed, the rest becomes far more facile:

- **Abstract the storage and data pools from the OS, storage types and applications**
- **Apply policies to the abstraction layer**

In simple terms, take a piece of data at the abstraction layer (i.e., it has already been written to some storage by a client or application). Then, classify it. Finally, apply a policy set to that data object, such as:

Value: Defined in financial or business value terms

Owner: Apply security policies based on access rights

Location: Storage Tiering and/or physical location

Retention: Inclusive of any deletion policies required

Through the use a common code base, it becomes easy to apply, for example, the following attributes to a data item:

- Security
- Encryption
- Classification
- Content Indexing
- Data Copy Services
- Policy Management
- Compression

The key differentiator is that these attributes are applied through policies and can therefore be automated. For customers using multiple backup, replication, reporting, and archiving tools this is of enormous importance. The operations team will be able to apply corporate or business policies against existing data sets without the need for modifying the storage, initial write policies and the plethora of tools designed to provide these functions currently. The business can be supported more transparently, response times are massively improved, and the cost of managing the storage function falls dramatically.

## 7. Conclusions

#### Main Findings:

- Tight integration is key
- Vendors building on a single software platform have the advantage

What has become clear is that multiple data management tools, wholesale data segregation, storage management silos and diverse processes have distracted many organisations from defining a concise and workable data and information management strategy. What is also apparent is that without the business collaborating with IT to define the value of information, it will be all but impossible for IT to deliver against an undefined requirement.

This point is critical for organisations. As more pressure is applied from external bodies in the way information needs to

be extracted from data and presented in an auditable fashion, it becomes imperative for organisations to reconsider their approach to the harvesting and mining of data and more so the capability to access its information resources “on demand”.

Much in the same way as organisations have consolidated storage and, more recently, server infrastructures, it is vital that consideration be given to data management tools and processes to not only deliver operational efficiencies but to also provide access to information and consequently knowledge through existing data sets.

**Unification is the key**

A storage architecture that provides tight integration between storage management, data management and data protection systems is key to effective end-to-end data lifecycle management. Unifying the traditionally separate functions of data movement with data management allows management of the entire storage stack from application to device as a cohesive whole.

In this area, starting from scratch has some advantages, as other players who gained much of their technologies through acquisitions are not so well positioned to provide integrated systems. Because of their legacy, these vendors are faced with the challenge of “rationalising” a number of disparate data management and analysis solutions together in order to deliver a cohesive, interactive platform.

The “data abstraction layer” appears to be the logical point at which to direct storage, Information and or data management policies. Removing the dependency of the data itself from the underlying physical assets provides a far more flexible approach to data and information mining, as well as to data migrations, backups, restores and business continuity

Data policies should be reflective of the business requirements, with the capability of being extended to incorporate new business drivers, for example in the manner in which worries around possible litigation have driven information retrieval requirements.

Overall, Quocirca recommends that an approach is found that combines abstraction with an integrated approach to data and information management, utilising in-built policy enforcement and security capabilities to ensure that the technical approach to data and information supports the dynamic changes in the business’s process needs.

## About CommVault

A singular vision – a belief in a better way to address current and future data management needs – guides CommVault in the development of Unified Data Management® solutions for high-performance data protection, universal availability and simplified management of data on complex storage networks.

CommVault's exclusive single-platform architecture gives companies unprecedented control over data growth, costs and risk.

CommVault's software was designed to work together seamlessly from the ground up, sharing a single code and common function set, to deliver superlative Data Protection, Archive, Replication, and Resource Management.

More companies every day join the thousands who have discovered the unparalleled efficiency, performance, reliability, and control only CommVault can offer.

Information about CommVault is available at [www.commvault.com](http://www.commvault.com).

CommVault's corporate headquarters is located in Oceanport, New Jersey in the United States. The EMEA Regional Headquarters is located in Reading in the UK, and the Oceania office is based in Sydney Australia.



## 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 in the following key areas:

- Business process evolution and enablement
- Enterprise solutions and integration
- Business intelligence and reporting
- Communications, collaboration and mobility
- Infrastructure and IT systems management
- Systems security and end-point 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
- Integrated print management

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, 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 at no cost. Quocirca's reach is great – through a network of media partners, Quocirca publishes its research to a possible audience measured in the millions.

Quocirca's independent culture and the real-world experience of Quocirca's analysts 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](http://www.quocirca.com).

**Contact:**

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

## Managing Customers and Content Maximising Revenue, Minimising Costs

### Contacts:

Elaine Axby  
Quocirca Ltd  
Tel +44 20 8874 7442  
[Elaine.Axby@quocirca.com](mailto:Elaine.Axby@quocirca.com)

Clive Longbottom  
Quocirca Ltd  
Tel+44 118 948 3360  
[Clive.Longbottom@quocirca.com](mailto:Clive.Longbottom@quocirca.com)

Fiona Moon  
CommVault  
Tel: +44 118 951 6500  
[Fmoon@commvault.com](mailto:Fmoon@commvault.com)

*The impact of the convergence of IT, Telecoms and Media is immense on all suppliers in the sector. Customers have an increasing number of information assets and can choose services from a range of suppliers. For companies wishing to succeed in the 'new world,' understanding and managing information is key. In order to achieve this, multiple content types and complex customer data must be able to be stored and retrieved quickly and efficiently.*

- **Convergence is finally here**  
Technology is enabling companies to deliver services outside of their traditional competences: mobile companies can deliver high speed Internet access; telcos TV services. Commercial pressures are forcing companies to broaden the range of services they offer, and they are seeking ways to generate incremental revenue from information services.
- **Customers can do so much more**  
This explosion of information possibilities offers choices to customers as never before; they can buy, exchange or store pictures, music, video, emails and messages which are generated from many sources and through a variety of companies in the telecoms and media space. Such organisations have significant value tied up in informational assets and need to be able to manage these for maximum commercial benefit.
- **Data remains in 'silos' – and the number and size of these silos is increasing**  
Traditionally, data sources were limited to basic call and customer data from network and customer management systems. As organisations increase the range of services offered, and more and more content is added, an explosion of data results. What is needed is a different approach to managing this data, pulling it together for the benefit of the business.
- **Information can be used for revenue generation and cost reduction**  
If information can be appropriately indexed and categorised, it can be used to better understand customers, their value and needs, to increase revenue by better customer retention and new service development. It can also help content providers better understand the value of their content and price it accordingly. It can help reduce costs by enabling more dynamic network provisioning and improving regulatory compliance.
- **A single 'abstraction layer' is required**  
We need to create an abstraction layer between the physical and logical assets and utilise this single resource pool as a single information library. This can be enabled through the use of virtualisation - at the storage and compute capability level. A service engine approach should identify and capture information at the point of creation, enabling the creation of composite applications for a range of business purposes.

### Conclusions

The challenges of managing an ever increasing amount of data in a converged telecoms, media and IT world are immense and increasing. Customer pressure and pressure from the business, will mean that companies will have to manage this data more efficiently, offering more targeted services and improving profitability. Customers will want content – be that emails, video, photos or music – where they want, on the device they want, at the time they want.

Creating a single abstraction layer has the potential to provide this data to service the business in a highly flexible and responsive manner to meet the challenges of this new world.

#### REPORT NOTE:

This article has been written independently by Quocirca Ltd to provide an overview of the issues facing organisations in telco and media around the management of intellectual property assets. 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.

An independent report by Quocirca Ltd.

[www.quocirca.com](http://www.quocirca.com)

Commissioned by CommVault

## 1. Introduction

The IT, telecommunications and media worlds are converging as never before. Changes in technology are enabling new services to be developed: mobile companies can use 3G spectrum to vastly improve the way in which customers access the Internet via their mobiles and are bidding for new slots in the spectrum which will enable the provision of TV over mobile phones. Increasing broadband speeds enable the fast provision of more and more bandwidth-hungry services. Digital TV enables the provision of broadband internet services, and fixed telecommunications providers can use next generation networks to transmit films and TV as well as simple email and Internet access.

The customer will have the ability not only to communicate via simple voice and text, or passively watch films or TV from a single provider, but also a vast choice of services to access an unparalleled choice about when, where and how to communicate and use information.

These developments have led to significant changes in the structure of the industry and the types of services offered. From a position where a single company offered a single type of service, be that fixed or mobile telecommunications, film, TV content or news, we now see service providers like Virgin Media offering “quadruple play” of broadband, Digital TV, fixed line and mobile phone services. Traditional telecommunications providers such as BT are now moving to offer a similar range of services, and mobile providers are typically offering fixed broadband as well as mobile services. Indeed, Tiscali has moved into the television space through a focused IPTV play. All depend on a limited number of content owners, who themselves are seeking to maximise return on their content assets.

Who will succeed in this marketplace is not clear: what is clear, however, is that this is a marketing game which needs a clear understanding of who your customers are, what they are doing and what they want to do. Companies need to be able to understand their customers and introduce – and withdraw if necessary – new services to profitable customers quickly and efficiently. They must also be capable of dealing with other content providers on a more dynamic basis – acting as an information pass through, enabling the content provider to gain a degree of branded access to the customer, while still leaving the customer ownership to the service provider.

Organisations in this sector have significant informational assets – customers and, in many cases, a growing portfolio of content. This paper explores the challenges of managing these assets to develop revenue generating services, cut costs and improve efficiency. It proposes an approach to creating a single logical layer of information assets which can then be accessed for a range of applications.

## 2. The challenges of managing information assets in telco and media

### Main Findings:

- Players in the converged market have a vast amount of data from a wide range of sources.
- They need to capitalise on these assets in order to develop new services, cut costs and improve returns.

For service providers, different types of customer information comes from a range of different sources:

- basic customer data: name, address, credit details, which services the customer is subscribed to. This data may include the type of equipment the customer has – mobile phone, type of PC, set top box, other device.
- customer history: service changes, upgrades, cancellations, interaction with customer services.
- call data: types of calls, time, duration, location if mobile – this may come from the network or billing engine.
- bandwidth used: how much bandwidth is being used and by whom.
- content: what the customer has used in the way of content; what was accessed immediately it was available; how many times accessed. How much was paid for; how much was looked at for free, but not taken further?
- contextuality – through what means is the customer accessing the data source? Does it make sense to provide it to them over that connection type, to that device and so on?

This data tends to sit in silos, as the business has grown: call data might come partly from the network, partly from the billing system, some CRM data from one system; that related to content services from another. This leads to challenges when trying to put together a coherent view of the customer base, understand which users are most valuable and develop services to maximise revenue generation.

In addition, many customers are looking for self-service, being able to change their services as and when they want to, and moving from the “standard” service bundles presently offered. Offering this type of service is difficult when data is so scattered. Customers also have more and more data which they need to manage: music, video, messages or emails which they want to keep. This will be stored in a number of places, both centrally and on local devices. Some services will also involve the streaming of live information to one or more people at the same time – without the capability for the recipient to store this stream, yet for them to be able to manipulate the stream as they use it, for example by pausing the stream.

Companies need to be able to use data to identify their most valuable customers and target new offerings at them. At the same time, many traditional markets are becoming saturated, and the ‘hangover’ from the boom years means sustained pressure to cut costs and improve returns. Companies have vital information assets, but can they use them wisely?

### 3. The need to use information for maximum gain

#### Main Findings:

- Data from disparate sources needs to be pulled together to be easily identified and used.
- Doing this can improve profitable customer retention and generate new revenue as well as cutting costs and improving the organisation's ability to meet legal and regulatory requirements.

A telecom or media organisation needs to be able to pull together data from the disparate sources previously referred to and have a coherent data source where single pieces of data are appropriately indexed and categorised so that they can be easily identified and contextually utilised in different situations.

The benefits of such an approach can be seen in a number of areas. The most important are in marketing and commercial negotiation:

- Marketing will be able to understand much better which services are being used, how much by whom, leading to improvements in customer segmentation, targeting of profitable customers and design of new services.
- New revenue generating opportunities can be created: some enterprise segments are interested in managed email services, some consumers will be interested in "if you liked that, try this"-type of offers. As the fixed and mobile, telecoms and media worlds converge, companies with 'quadruple play' offers will be able to offer customers a managed content service – the content on the device they want, where they want, at the time they want.
- The value of content in general will be better understood – owners and content providers can see what is accessed when and by whom, and can price wholesale or retail prices accordingly
- Customers will be better able to navigate information repositories themselves, utilising meta data tags to find similar items such as films and TV programmes of a similar nature, so lowering service and provisioning costs.

Other departments will also benefit from better information management:

- Network efficiency can be improved: for example, greater intelligence on what types of information might cause a surge in bandwidth demand can help service providers apply existing tools more effectively to dimension and provision the network on a more dynamic basis, so optimising response times and maintaining customer satisfaction.
- Regulatory compliance will be improved: data retention requirements in Europe have increased as a result of the EU Data Retention Directive adopted in 2006 and being implemented now. This potentially adds significant costs to service providers who now have to keep subscriber and call data, device and location data in a form searchable by the relevant authorities. The regulatory environment is likely to evolve in the coming

years, and the organisation will need to be able to respond quickly and cost-effectively to new legal and regulatory demands.

To really maximise the benefit it gets out of its data assets, the organisation needs to take a different approach to its data management. Below, we suggest an approach creating a single abstraction layer, which allows a horizontal approach to be taken to data, enabling applications to be built on that layer as required.

### 4. Creating a single abstraction layer

#### Main Findings:

- An abstraction layer created through virtualisation between the physical and logical assets enables the creation of flexible data management.
- Such an approach should be based around an engine that is there as a service itself, enabling the creation of a range of composite applications and better managing existing storage assets.

To create a flexible environment where all data assets are more easily managed, the way to go is to create an abstraction layer between the physical and the logical assets. At the physical level, there are the various parts of the network, the servers and the storage. At the logical level, pools of resource are created that are callable by any other logical entity – whether this is an existing application, a web service or whatever. Commonly viewed as "virtualisation", the overall effect of a coherent approach is to divorce the dependencies of the business process and application needs from the physical constraints and problems found at the hardware level.

As an example, at the storage level, there may well be multiple different storage assets, with different classes of disk drives (high speed and/or high volume). If we want to use these assets effectively without using virtualisation, we run in to problems around how we allocate the storage at a "one to one" level between physical applications on physical servers to physical storage. If an application creates more output than expected, the need to extend the storage becomes an issue, and it may be necessary to replace the existing storage asset to replace it with a larger system with sufficient headroom for the future. During the time required to carry out this upgrade, application availability will be hit – and this may well impact the capability of the organisation to generate revenues or maintain customer service.

Now, if this hard linking is replaced with the use of an abstraction layer against a virtualised storage environment, we can create logical units of storage against the physical assets. For example, all high speed storage can be put together as a logical "tier 1" group, which we can then allocate as required to our most demanding applications, such as video on demand (VOD). Email storage can be allocated to our lower performant "tier 2" assets. Now, should we need extra storage for any of our applications, we can add further physical assets that will be embraced within the existing environment and will automatically join the pool of resource.

Therefore, this approach can provide not only greater flexibility in how the underlying physical assets are used,

leading to greater utilisation levels, but also maximises service uptimes and provides a means to move to newer, faster storage technologies as they emerge without looking to mass replacements. The existing “tier 1” storage assets can simply be reallocated as new “tier 2” resources.

Again, if information is stored against physical data stores, it needs to have “hard” physical pointers to it, and these pointers have to be changed as the information is moved about – either from one main informational store to another, or as we age the data from tier one storage to tier two and so on. If there is a single logical pool of information, these pointers are maintained automatically – and we can therefore define information policies that align far more closely with the customer and the provider’s needs.

Also, through identifying and tagging information at the point of creation in such a way so that we can interrogate the whole environment in an open yet secure manner, we find that we create a highly flexible and responsive environment where compliance is built in, rather than built on.

Taking such an approach to the abstraction of data and information from the physical layer creates a solid platform for the future. Information providers can be more easily brought on board, and agreements on information management and digital rights management (DRM) are far easier to agree and manage.

This approach dictates that the underpinnings be based around an engine that is there as a service itself, an engine that comprehensively deals with information on an end-to-end basis. This engine should be able to identify and capture information at the point of creation, should be able to apply security and tags based against a corporate policy, must be able to index the information to make identification easy, and must be highly responsive to the needs of the business itself.

Through the use of such an informational engine, we can then start to build our compliance oriented architecture. As the underlying information is now available as a service, it is possible to create composite applications that can be utilised where a user can easily create self-service requests that are self-fulfilling, where information providers can automatically plug in new offerings as feeds to the system, where information assets age effectively across different physical storage types and where the service provider can make dynamic offers directly to its customers – and even to prospects on a far more one-to-one basis.

## **5. Conclusions**

The challenges of managing an ever increasing amount of data in a converged telecoms, media and IT world are immense and increasing. Customers are going to demand ever more personalised services and many will want to manage their own information services. They will want content – be that TV, video, email, photos or music - on the device of their choice, at the time of their choosing and wherever they are. This customer pressure will mean that companies will have to manage their customer and data assets more efficiently, offering more targeted services and improving profitability.

Pressure on the business to cut costs and operate more efficiently means that existing storage assets will need to be managed rather than replaced, and the data from existing infrastructure assets used to improve service levels and cut costs.

Creating a single abstraction layer for data enables it to be managed far more easily and be available quickly and easily for a wide range of applications. Enabling data to be called as a service creates a far more responsive and flexible platform for the business to work with – and to meet the upcoming and dynamic challenges of the markets.

## About CommVault

A singular vision – a belief in a better way to address current and future data management needs – guides CommVault in the development of Unified Data Management® solutions for high-performance data protection, universal availability and simplified management of data on complex storage networks.

CommVault's exclusive single-platform architecture gives companies unprecedented control over data growth, costs and risk.

CommVault's software was designed to work together seamlessly from the ground up, sharing a single code and common function set, to deliver superlative Data Protection, Archive, Replication, and Resource Management.

More companies every day join the thousands who have discovered the unparalleled efficiency, performance, reliability, and control only CommVault can offer.

Information about CommVault is available at [www.commvault.com](http://www.commvault.com).

CommVault's corporate headquarters is located in Oceanport, New Jersey in the United States. The EMEA Regional Headquarters is located in Reading in the UK, and the Oceania office is based in Sydney Australia.



## 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 in the following key areas:

- Business process evolution and enablement
- Enterprise solutions and integration
- Business intelligence and reporting
- Communications, collaboration and mobility
- Infrastructure and IT systems management
- Systems security and end-point 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
- Integrated print management

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, 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 at no cost. Quocirca's reach is great – through a network of media partners, Quocirca publishes its research to a possible audience measured in the millions.

Quocirca's independent culture and the real-world experience of Quocirca's analysts 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](http://www.quocirca.com).

**Contact:**

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