1

Introducing Utility Computing

In the 1970s and 1980s, mainframe computing comprised huge global computing systems. It was expensive and had a pretty bleak user interface: but, it worked. In the early 1990s enterprises moved to highly distributed client/server computing which allowed IT to deploy PC client systems with, on the face of it, lower cost and a much better end user experience. By the late 1990s, Internet computing allowed systems with the mainframe's centralised deployment and management but with rich PC-like browser-based user experiences.

Now, the industry is in that age of *Utility Computing*. Utility Computing is a term the IT community has adopted that represents the future strategy of IT. No vendor is embarking alone in this approach – all the major vendors have their own version of this vision. But whatever it is called, Utility Computing represents an evolution of the way corporations use IT. So, what's different about Utility Computing?

Utility Computing is the first computing model that is not just technology for technology's sake; it is about aligning IT resources with its customers – the business. Shared resources decrease hardware and management costs and, most importantly, enables charge back to business units. Utility Computing also has autonomic or self-healing technologies, which comprise key tools for the CIO to make business units more efficient. But it isn't possible to buy

Data Lifecycles: Managing Data for Strategic Advantage Roger Reid, Gareth Fraser-King and W. David Schwaderer

Utility Computing off the shelf because Utility Computing will evolve over the next 5 to 10 years as technology advances. Organisations, however, can help themselves by setting up the correct building blocks that will help intercept the future. Most enterprises now use available products for backup and recovery. Large organisations can also provide numerous IT management functions as a utility to the business.

If parts of a business are charged back for IT services, then the size of that charge back becomes a key measure of success. Data storage, for example, has costs associated with it the same way that paper-based filing cabinets, clerks, floor space and heating overheads did 20 years ago. Keep in mind that these solutions must provide a framework across heterogeneous IT infrastructures that provides IT with the ability to manage and justify all assets back to the business, as well as provide the business with continuous availability of mission critical applications and data. Even if the organisation decides not to bill back, the insights can prove immensely valuable.

Attempting to make realistic IT investment decisions poses a dilemma for business leaders. On one hand, automating business processes using sophisticated technology can lead to lower operating costs, greater competitive advantage, and the flexibility to adjust quickly to new market opportunities. On the other hand, IT spending could be viewed the traditional way – a mystery, essentially due to the view of IT as an operational expense, variable cost, and diminishing asset on the corporate balance sheet.

By treating IT as an operation, organisations combine the costs, making it next to impossible to account for individual business usage. From an operational perspective, this means that not only are usage costs hidden in expense line items, but also the line of business has no way of conveying its fluctuating IT requirements back to the IT department. Moreover, this usually leads to the IT department having a total lack of understanding for the business requirement for service levels, performance, availability, costs, resource, etc. Hence, the relationship between IT spending and business success is murky, and often mysterious. So Utility Computing attempts to simplify and justify IT costs and service to the business.

Utility Computing effectively makes IT transparent. In other words, a business can see where its funds go, who's spending the largest funds and where there is wastage or redundancy. Utility

Computing means that lines of business can request technology and service packages that fit individual business requirements and match them against real costs. This model, then, enables a business to understand IT purchases better, together with service level choices that depend on the IT investment. When making IT purchasing decisions, historically businesses arbitrarily threw money at the IT department to 'do computing' to make the system more effective. Now, Utility Computing enables businesses to obtain Service Level Agreements (SLAs) from IT that suit the business.

Transparency of costs and IT usage also enables organisations to assess the actual costs associated with operational departments. In the past, this was not possible because IT was simply seen as a single cost centre line item. Now, IT can show which costs are associated with which department – how much storage and how many applications the department is using, the technology required to ensure server and application availability, together with how much computing power it takes to ensure that IT provides the correct level of service. This visibility allows IT departments to understand storage utilisation, application usage and usage trends. This further enables IT departments to make intelligent consolidation decisions and move technological resources to where they are actually needed.

Giving IT the ability to provide applications and computing power to the business when and where it is needed is essential to the development and, indeed, survival of IT. By being able to fine tune IT resources to meet business requirements is essential in reducing overall cost and wasted resource. It saves time and personnel overheads. Not only does it mean the end user experience is dramatically enhanced, but also the visibility of how IT provides business benefits becomes apparent. We may characterise IT as a utility, but what we really mean is providing IT services when and where they are necessary; delivering applications, storage and security, enhancing availability and performance, based on the changing demands of the business and showing costs on the basis of the use of the IT services provided.

The Utility Computing approach not only provides benefits to the business but also to the IT department itself. As IT begins to understand the usage from each of the business units, IT then has the ability to control costs and assets by allocating them to specific business departments and gives IT management a better

understanding on how IT investment relates to the success of business tasks and projects. The utility approach gives IT the ability to build a flexible architecture that scales with the business.

The challenge for many IT departments is deciding how best to migrate current IT assets into a service model which is more centralised, better managed, and most importantly, better-aligned with the needs, desires and budgets of departmental users. This means increasing servers and storage utilisation through redundancy elimination.

Utility Computing methodology can provide significant cost savings. By delivering IT infrastructure storage as a utility, organisations can:

- reduce hardware capital expenditures;
- reduce operating costs;
- allow IT to align its resources with business initiatives;
- shorten the time to deploy new or additional resources to users.

Provisioning enterprise storage – including storage-related services such as backup and recovery and replication – within a service model delivers benefits for IT and storage end users. It can maximise advantages of multi-vendor storage pool resources, improve capacity utilisation, and give corporate storage buyers greater leverage when negotiating with individual vendors. This service-based approach also allows storage management to centralise, improving administration efficiencies, allowing best practices to be applied uniformly across all resources, and increasing the scope for automation.

A storage utility delivers storage and data protection services to end users based on Quality of Storage Service (QOSS) parameters of the service purchased. Delivery is automatic. The end user need not know any storage and network infrastructure nuances to utilise capacity allocations or be assured of data protection. At the end of each month, billing reports detail how much storage each consumer used, the level of data protection chosen, and the total cost. This allows each consumer to assess storage resource usage – whether it is physical disk allocations or services offered to secure the allocations – and make decisions about how they plan to utilise the resources in the future.

A storage utility strengthens the IT department's ability to satisfy end user service level demands. By clearly stating the

expected service levels of each packaged storage product, the IT department helps end users accurately map application needs to storage-product offerings. This gives the IT department a clear understanding of the service-level expectations of business applications. End users of the business application benefit by knowing that IT is able to live up to the service level it has defined.

Just as a storage utility can use storage management software and Network Attached Storage (NAS) or Storage Area Network(s) (SAN) technologies, a server utility can similarly 'pool resources' and automate rapid server deployment for specific critical applications to meet specific business requirements.

Automating application, server, and storage provisioning, as well as problem management and problem solving through policy-based tools that learn from previous problems solved, will play a large part in future advances in deploying utility storage. Predictions of future usage, as well as automated discovery of new applications, users, devices, and network elements, will further reduce the IT utility management burdens as it evolves from storage to other areas.

1.1 Real problems and real solutions

1.1.1 Real issues identified – regulation, legislation and the law

Regulations traditionally dealt with business information management via paper-based audit trails. But these regulations have become redundant over the years – no paper, no paper-based audit trails to follow. Legislation needed a decent make-over. It took a while, but regulations have now begun to catch up with the movement of data from paper-based storage to electronic data storage devices. To exacerbate matters on the regulatory front, we have recently seen terrorist acts and corporate scandals that have increased the amounts of data that organisations have to store. The effect of these additional regulations is to exponentially increase the amounts of data that organisations have to store and for longer periods.

Now, generally storage is relatively cheap, however, the issue is not the storage of the data so much as the retrieval of the data.

Because there is so much data being saved it is much like looking for the proverbial needle in the haystack. Organisations, therefore, must have the ability to understand the relative importance of their data within its lifecycle as well as have ways to find it in an open system that historically has had no due process behind its filing methodology.

So, storing information effectively is unquestionably vital for organisations, but with data volumes rising frighteningly and a growing need to make archived data available both for end users and to comply with legislation, the way IT departments approach storage is critical. Although the storage price per gigabyte may be dropping, simply installing new devices is not always a perfect solution. Rather than making data harder to retrieve and contributing to rising costs for support and maintenance, many organisations are looking to reduce the complexity, inefficiency and inflexibility of their data centre environments.

And so *Data Lifecycle Management* (DLM) was born. Previously, Hierarchical Storage Management (HSM) existed simply so that an organisation did not store old data on its most expensive disk. Now DLM has become the 'hot' subject. How do we manage data and retrieve it at will? Well, simplistically you could tag the data and then use a decent search engine.

Actually, it hasn't taken organisations long to work out that, not only do they want to be able to retrieve data but also to store it logically so that like files are stored in the same place – hence, *Information Lifecycle Management* (ILM). ILM in itself suggests some due process or implied activity that has occurred to the 'data'. This is where technology is searching for a utopian solution.

Total Lifecycle Management (TLM) is the technology that will make all and/or any document(s) retrievable in an instance; the data is logically stored on the most appropriate medium for the correct length of time and then deleted from disk or the tape destroyed at the right time – automatically.

1.1.2 More regulation, legislation and the law

Failure to retrieve data becomes increasingly critical to organisations when new regulations require data retrieval, an audit trail proven, as well as the ability to prove originality and what has

happened to the data when, where, how, and by whom. There are many examples of companys' prosecutions and fines, although there is a lack of high profile prosecutions simply because organisations try to play down any large fines because of the potential bad publicity.

The UK Information Commissioner's Annual Report lists prosecutions in the 12 months between 1st April of the previous year and 31st March of the year of its annual report. In the last report, there were 10 defendants convicted – in all of these cases the defendants were convicted of multiple breaches of the Data Protection Act (UK) with fines up to £5000. (Potentially fines can be up to £5000 in the magistrates court and unlimited in the Crown Court.) Prosecutions have recently been approached on a 'per data subject' basis, i.e. where a company has breached the Data Protection Act (UK) in respect of one individual a conviction has been sought and a fine imposed; where the company has breached the Data Protection Act (UK) in respect of a number of individuals a conviction has been sought and a fine imposed in relation to each individual. Therefore, according to this approach, where the personal data of 500 data subjects has been misused, 500 fines of, say, £5000 could be imposed (£2,500,000 or $4,000,000 US).

And not only is there new legislation to deal with the new phenomenon of electronic data, but old laws are catching up. We now have of examples of entertainment exploiting large enterprise organisations who have no idea what they are storing in their vast data warehouses. In fact, most third-party or copyright infringements relate to the sharing of electronic entertainment media. DVDs and CDs have made third-party infringement a big issue. A recent news report indicated that a media company, which determined that music piracy was on the increase, decided to look at, not the cause of the copyright theft, but the holding company . . . so to speak.

Previously, someone taping a vinyl record was a nuisance, but now with perfect reproductions possible with each copy, copyright infringement has become a big problem. Peer-to-peer music sharing may well be neat technology, but unfortunately it's illegal to actually do any sharing unless you both own the rights to the music (if that was the case why bother sharing?). But suing an individual for breech of copyright is hardly worth the bother. Now consider an employee putting their own music onto their work computer, no problem so far. Suppose these guys are members

of the Musicians' Union and so the last thing they are going to do is share the music – which they know is illegal. So, are they OK? No. . . .

What happens when their workstation or laptop is backed up? All the MP3 files back up onto an organisation's network server and then migrate onto offsite storage tapes. Before you know it, you have multiple illegal copies of redundant data, all illegal. To make your day even worse, not only are you storing illegal redundant files on valuable disk space but the media company the music belongs to in the first place can then take you to court for big monetary fines.

Recent Forrester research revealed that 2/3 of all organisations in the USA in 2003 had illegal music files held on their servers. Not only are they storing something illegal, they don't really want to store it in the first place. Typically, in most organisations, 30% of all stored data is illegal or simply rubbish. This, of course, has a storage management and media cost impact. It also has an immediate and recurring impact on the time it takes to backup data. Eliminating this data thereby helps reduce the data growth rate. All these considerations are of vital importance to organisations over the next few years.

1.1.3 Current storage growth

Finally, data is quite rightly viewed as a key aspect of an organisation's operation and success. To underline the fact that data is one of an organisation's most important assets, consider that managing information badly through inept retrieval or illegally held data can have enormous financial implications. The sheer volume of digital information is increasing exponentially. Web sales, email contracts, e-business systems, data demanding sales, marketing and operational systems – all of which are the lifeblood of most modern organisations – not to mention managing wireless, and remote and handheld devices, together with multimedia usage, all lead to heavier data traffic and more storage requirements, with larger and more files being saved.

All this stuff needs to be saved, stored, retrieved, monitored, verified, audited and destroyed, not just so the organisation can do business, but also to comply with data retention legislation, just

so the organisation can continue commerce without the threat of financial penalty or operating licence withdrawal.

1.2 New storage management

Organisations need a new way to manage storage. The IT world has turned their eyes towards DLM/ILM/TLM. The concept of Data Lifecycle Management/Information Lifecycle Management provides IT organisations with a better way to manage a wide variety of data or information, this includes traditional structured files, unstructured data, digital media (sound, video and picture files) and dynamic web content. DLM/ILM will index all types of content and logically store content with like type content on the most appropriate storage media for that content type, within its lifecycle. This helps organisations improve access, performance, utilisation, and costs, to ensure compliance as well as providing customers with an efficient service.

Many vendors are still advancing HSM tools and trying to characterise them as ILM solutions. However, DLM or ILM solutions are generally not successful if they have been developed from legacy HSM technology. Even if a magnificent tool appears to do everything asked of it, IT departments must still understand that the building blocks required for a successful ILM strategy are in the storage management layer and the long term efficient management of information throughout its entire lifecycle.

A DLM/ILM strategy cannot, and must not, be undertaken solely by the IT department: that would be an impossible task. How could IT possibly know what policies to build for which regulation, and what business requirement is needed to ensure the service they provide to the business is accurately documented and delivered?

In many cases regulatory compliance is simply thrown over the proverbial wall at the IT department, because 15–20 years ago IT made those decisions. This is ill-advised. Currently, numerous projects are occurring to satisfy legislative appetites. Soon, organisations will realise that getting compliant at any cost is simply infeasible and far too costly. This conundrum is an ongoing process that will continue to change and evolve from day to day.

It makes much more sense to build a storage infrastructure based on one of the more widely known quality standards, ISO

or ITIL for example. This helps prevent having a large number of simultaneous projects that potentially contradict each other. The infrastructure then supports *policy based management* – here, the business makes policy and IT implements it. To make things easier tools already exist on the market that can extract data on most platforms and at a granular level, as well as provide version information and dynamic rule application to data upon creation that intelligently travels with the data throughout its lifetime until it is eventually destroyed.

Already, technologies are appearing that can unify the management of ILM policy setting as well as view the whole storage environment from servers to an offsite archive. These tools, adjuncts to traditional storage management software that has evolved from the mid 1990s, provide a link between the storage layer and the server and end user. They effectively give business managers visibility into the way their legacy data is stored to ensure that it is being done with the most relevant protection, availability, and compliance requirements, over its entire lifetime.

The bottom line is that all organisations need a robust, scalable record retention and retrieval strategy. They need to store all their data in a secure location, that is cost effective and efficient, for as long as is necessary, that is resilient over time, and compatible with legacy and future media formats and technologies. Data must be stored for fixed periods of time (sometimes as long as 90 years) and, in some cases, on a storage medium with specific properties such as WORM. During an audit, organisations also require the ability to discover and retrieve electronic records in a timely manner. Therefore, efficient access to information and consistent availability is also necessary. To be effective, organisations need to be able to produce requested data in a timely manner, which can often mean within as little as 48 hours, or risk a more in-depth audit or worse.

Organisations also need to be able to guarantee data integrity to protect against alteration and be able to verify originality; in other words to ensure that the data is original and has not been altered in any way which, of course, includes newer application versions. And organisations need to be able to store original content, unalterable media, as well as new 'unstructured' file types including: memoranda, email, Instant Messaging, and other forms of digital information.

With many organisational processes moving from paper-based operations, compliance regulations require companies to demonstrate internal controls and processes in order to document what they do and how they do it, as well as demonstrate adherence to the regulations, so that in the event of an audit, they can show who had access to the data, when, and what actions were performed. A system failure or lack of visibility into the system is not an excuse for noncompliance.

Therefore, information must always be available for review by an auditor, with efficient accessibility of information and consistent availability – and this also requires the ability to produce reports that reflect origin of data and activity in real time.

1.2.1 What are the things organisations need to consider?

So the challenge for the IT manager remains: in the business world organisations need to deal with ever increasing volumes of information that are ever diverse, and increasing in size with every release of Microsoft Office. And, IT has to do this without extra budget or strain on the IT department's workforce, who, incidentally, are already working a 65-hour week. In addition, IT can't try to manage the storage resources by simply just adding more inefficient direct-attached storage devices, because that just doesn't work in the long term – and how do you successfully manage disparate storage devices anyway?

If storage growth is compounding at 50–100 % year, an organisation with one terabyte this year will potentially reach 32 terabytes in three years including backups. So not only have you got to put this stuff somewhere (more storage), but then you have to manage as well. Users still expect instant uninterrupted data; administrators face increased scalability and performance requirements, which are both initially unmanageable and invisible, with restricted (perhaps decreasing) budgets. In addition, board level executives need to ensure their company information is protected, accessible, and retained according to the latest worldwide, international, country and local regulations – regulations that are in their 10's of 1,000's and are constantly changing.

As a single example, in the USA alone there are over 10,000 US Federal Regulations surrounding electronic information retention. Extraction of data archive point-in-time views are becoming normal.

1.2.1.1 The problems

The main problems for an organisation are as follows:

- static or decreasing IT storage management budgets;
- multi-platform skills shortage;
- fewer IT system Admin Engineers;
- more sites, more data, more systems – all needing management;
- unstoppable data volume growth;
- globalisation – organisations now need to be available $24 \times 7 \times$ forever;
- compressed to zero backup windows;
- increased regulatory legislation around data management, IT and corporate governance;
- new communication types that need some sort of business policies set against them in risk mitigation
- inability to manage or control storage costs.

1.2.1.2 Things to consider

The main things an organisation needs to consider are as follows:

- What types of data does your organisation hold?
- Which of these data types need to be held?
- For what length of time does this data need to be held?
- Is any of this data likely to be used in the future?
- How critical is the data to the business?
- Who needs to access it?
- How quickly do they need to access it?
- Does it need to be held and produced in its original state (WORM)?
- If required, could you deliver every single instance of one type of specific data required by government legislation?

1.2.2 What does data lifecycle management mean?

1.2.2.1 What is IT Lifecycle Management? (Defining DLM/ILM/TLM)

Data Lifecycle Management (DLM) is a policy-based approach to managing the flow of an information system's data throughout its lifecycle – from creation and initial storage, to the time it becomes obsolete and is deleted, or is forced to be deleted through legislation.

DLM products attempt to automate processes involved, typically organising data into separate tiers according to specified policies, and automating data migration from one tier to another based on those criteria. As a rule DLM stores newer data, and data that must be accessed more frequently, on faster, but more expensive storage media. Less critical data is stored on cheaper, but slower media.

Early types of DLM tools included HSM. The hierarchy represents different types of storage media, such as RAID (redundant array of independent disks) systems, optical storage, or tape, each type representing a different level of cost and speed of retrieval when access is needed. Using an HSM product, an administrator can establish and make policies for how often different kinds of files are to be copied to a backup storage device. Once the guideline is established, the HSM software manages everything automatically. Typically, HSM applications migrate data based on the length of time elapsed since it was last accessed, whereas DLM applications enable policies based on more complex criteria.

The terms Data Lifecycle Management (DLM), Information Lifecycle Management (ILM) and Total Lifecycle Management (TLM) are sometimes used interchangeably. However, a distinction can be made between the three.

- DLM products deal with general file attributes, such as file type, size, and age.
- ILM products have more complex capabilities. For example, a DLM product allows searching of stored data for a certain file type of a certain age, whereas an ILM product allows searching of various types of stored files for instances of a specific piece of data, such as a customer number.
- TLM products allow formulating complex requests across multiple storage tiers and heterogeneous operating systems to

provide a more complete approach to managing all structured and unstructured data.

Data management has become increasingly important as businesses face compliance consequent to modern legislation, such as Basel II and the Sarbanes-Oxley Act, which regulate how organisations must deal with particular types of data. Data management experts stress that DLM is not simply a product, but a comprehensive approach to managing organisational data, involving procedures and practices as well as applications. Fundamentally what has happened over the last 15 years, since the advent of 'Open Systems', is that the ability to process information in a coherent, cohesive and consistent manner has been lost, or at the very least, seriously mislaid.

It would be quite a powerful technology that could examine an organisation's data storage and re-file all data consistently, in an intelligent manner, and that would allow the organisation not just to retrieve information easily (because every file system would be standard), but to store that data logically together in appropriate batches with like-times for deletion, as well as migrating data to upgrade storage – keeping the integrity of the data intact – and bringing a copy of the appropriate version software with it so it can be read in the future.

Suppose it is important to find data in the future and that it is not conveniently located where one would expect to find it . . . necessarily. So what? What drives the need for DLM or ILM products and services?

- Emerging regulatory and compliance issues (Data Protection, HIPAA, International Accounting Standards, Sarbanes-Oxley, Basel II, etc.), which drives
  - unbridled data growth (both structured and unstructured data), which promotes
  - the variability in value of data that an organisation owns.
- Organisations continue to pressure CIOs to manage more with less, and to control costs, so
  - it is becoming increasingly difficult, nay, down right impossible, to manage an organisation's data manually across an increasingly distributed and complex environment with any kind of hope of success.

This has not always been the case. Back in the 1970s and 1980s, mainframes kept all the data in logical file systems as previously mentioned. However, since the arrival of Client Server/Open Systems in the early 1990s, the art of information management has been lost. Basically, personal record keeping has become a chaotic free-for-all. Each individual stores and saves his or her data in different ways. According to a leading analyst, 60 % of all data is unstructured – our email, file and print servers (word docs, XLS spreadsheets etc.). How then, does one find a specific piece of data from an employee who worked at the company for four years and left two years ago? With no 'due process' it's not easy – and there have been several organisations that have been billed in the six figure region just to find and retrieve the data.

To make matters worse, all offices started to go 'paper-free' around the early 1990s. Prior to this, all organisations had to store their information in hard copy storage systems, including Microfiche, all of which were fairly sophisticated with offsite, fire-proof, storage facilities and processes behind filing and record keeping as well as audit trails to show due diligence. However, since the early 1990s these hard copy storage warehouses have slowly but surely disappeared, replaced with electronic data warehouses. Paper records have disappeared and have been replaced with electronic data. To make things even worse, organisations now have additional communication methods and a range of electronic routes to market.

1.2.3 Why is IT lifecycle management important?

It should be obvious that organisations must manage and store data more effectively. The upside is that ILM/DLM makes good business sense: in fact, that's why it existed in the first place. ILM/DLM is a prerequisite for good corporate governance, but is also an integral part of good business conduct. It protects reputations and manages risk, as well as promoting a safe, secured transaction environment. It protects global financial market safety and stability as well as tracking suspicious customers' movement. It adds value to customers' confidence and with it competitive advantage. It helps prevent terrorist money-laundering activities and harmonises international regulatory approaches. Why wouldn't anyone want to know about Data Lifecycle Management?

1.2.4 Goals of data lifecycle management

Data is one of the most important organisational assets.

The above statement must be pure plagiarism. How many books, white papers, web sites or articles have made that statement? How many analysts, journalists, sales managers, business managers, business gurus, marketing managers, operational managers, database administrators and system administrators and have been bleating on about the benefits of looking after an organisation's data and information? Surely businesses must have caught on by now? Possibly, but unfortunately, probably not; but the scene is changing. Instant data gratification is out and data longevity is in. With an increasingly compliant and litigious society, data must be kept and accessed for longer.

Data as an asset is important in providing organisations with valuable information. Data becomes information and information becomes knowledge. This book discusses the differences between Data Lifecycle Management, Information Lifecycle Management, and Total Lifecycle Management in detail and examines the dichotomy at length. Although the principles behind the three concepts remain fundamentally different, it is all still data. Information management suggests that an organisation has done something intelligent with its data, and knowledge suggests that some cognitive process has been applied to that information.

From a technological point of view it is easy just to refer to DLM, and so initially we need to describe the fundamental goals of Data Lifecycle Management and its platform.

- **To make an organisation's data accessible.** All data should be readily available to support the businesses to which it is purposed. Availability requirements should not be restrictive.
- **To have an adaptable design and architecture.** Data continually changes. Hence, the processes, methodologies and underlying technologies that manage it should adapt to meet growing data demands.
- **To provide operational security to the asset.** The data management platform coupled with its process and methodology should provide auditing, tracking, and controlling mechanisms to manage the data effectively. Specifically, it must provide a complete management infrastructure that affords greater visibility into its daily use.

1.2.4.1 What are the technology trends we will see over the next few years?

Here are some of the expected technology trends:

- **Hierarchal Storage Management** (1997–2005). Vendors are already veering from using the HSM term and using both DLM and ILM instead. Although the term ILM is fraught with confusion and conflicting interpretations from various vendors (depending what technology they offer), vendors are already introducing numerous technologies that will be the starting point for the development of automated ILM products.
- **Data Lifecycle Management** (2003–2006). In the last few years, DLM products have emerged – the father of ILM if you like. With the increase in retrieval requirements through compliance issues and other emerging regulations, organisations have started exploring new ways to backup, store, manage and track their most critical data starting with virtual tape and disk-based backup. They have also started to implement some basic tiered storage capabilities, such as moving 'stale' data from their high-performance disk arrays to more cost-effective systems or deleting irrelevant data (MP3 files) altogether.
- **Manual Information Lifecycle Management** (2004–2007). This technology has the capability to index, migrate and retrieve data as well as prove its authenticity – on any part of the infrastructure. Although still a manual process when setting policies against business requirements, this is the point where ILM gives organisations the ability and technological intelligence to implement more-powerful storage policies. Some organisations will utilise virtualisation applications which logically group many arrays into a single 'virtual' storage pool and host them on emerging 'smart' storage switches. Manual ILM will provide users with a single logical file system view that is, in reality, scattered across multiple media types in multiple locations. ILM will enable companies to move data fluidly within the storage infrastructure as their evolving policies dictate, while shielding administrators and users from the underlying complexity.
- **Automated Information Lifecycle Management** (2006–2008). Automated ILM will integrate products that manage storage, virtualisation, and the data itself. Numerous storage management stack aspects will be imbedded into the infrastructure.

Compliance capabilities (retention, deletion, etc.) will also be imbedded into a number of storage management products and be well established in some vertical markets, especially the financial sector. The transition from manual ILM to automated ILM will require additional technologies in order to manage data as 'information', as opposed to managing it as 'data'.

- **Automated Total Lifecycle Management** (2007–2010). The total cost and value of a piece or set of data depends on every phase of its lifecycle, as well as on the business and IT environments in which it exists. TLM will automate the way that organisations look at their entire data set. TLM will offer organisations the ability to protect against media obsolescence, legacy data, future hardware changes, as well as dealing with all manner of diverse mobile assets, automatically managing storage costs and data movement (to lower cost storage options) as required, providing audits where required without the need for manual intervention. In other words all an organisation needs to do is decide on the data policy and TLM does the rest.