IBM Global Technology Services
September 2008

Trouble in store

Data storage isn't working – is it time to impose the class system?

As companies, government departments and other organisations pile up data ever more quickly, they face growing costs and inefficiencies. The global air transport industry has its own struggles with congestion, and may have a valuable lesson to offer.

Contents

Networks under pressure

The worlds of information technology and air transport are intimately linked. Airliners are no longer designed on drawing boards but on screens, in deep three-dimensional detail. They depend on computers for navigation, management of their onboard systems, even their very safety and stability as they cruise the skies at 600mph. Controllers on the ground are backed up by computer decision support as they manage heavily used airspace. And the airlines rely on software systems for dozens of functions – receiving bookings, issuing boarding passes, maximising their earnings, scheduling their crews, and many more.

But they also have something else in common, and it's much less desirable. Both IT and air transport are beginning to suffer the problems of excess. The airlines and their passengers suffer growing delays as the air traffic management system struggles to fit more and more flights into the available airspace and runway capacity. Costs mount as aircraft queue on the ground with their engines running and circle in holding patterns as they wait their turn to land. And the planet labours under the burden of millions of tons of extra CO_2 emissions.

Excess manifests itself differently in IT but with many of the same results – delays, more cost, more damage to the environment. The culprit is data – individuals, companies and other organisations are producing electronic information at a rate that's increasing exponentially, stressing the ability of the IT industry to give them what they need to store and retrieve it efficiently.

An extraordinary turning point was reached in 2006. That was the year when electronics manufacturers for the first time produced more transistors than the world's farmers grew grains of rice – at least 26,000 million of them. Packed on to the chips that power personal computers, smartphones and a host of other devices, the fundamental building blocks of information technology each cost about the same as one printed newspaper character.

Highlights

Moore's Law predicted soaring computer performance – but not the data tsunami that followed

These figures are the living proof of Moore's Law – the proposition, first advanced in 1965, that each new design of chip could be expected to do twice as much as its predecessor, leading to an exponential rise in performance matched by a corresponding fall in the cost of computing power. But there's a downside, and that's the data tsunami now engulfing businesses and governments.

It brings with it soaring demand for storage – a recent assessment put the growth in the market at 50 per cent a year for the discs, networks, management tools and related services needed to store information reliably and securely. What's more, that figure is itself growing. While this may be good news for some suppliers, it spells trouble for user organisations who are increasingly concerned about the resulting costs, inefficiencies and energy consumption. It also represents a challenge for those suppliers who see themselves as partners to their clients and are committed to helping them to operate efficiently, profitably and in an environmentally responsible way.

The survival of an enterprise could depend on its ability to master the data storage problem

Ultimately, the very survival of an enterprise can be at stake. In an economic environment where there is increasingly no room for slack on costs, quality and speed to market, getting on top of the data problem can make all the difference between success and going out of business.

Highlights

Accumulated data affects a business the way weight and drag affect an aircraft, weighing it down and increasing its operating costs

Weight, drag and cost

The goal of every aircraft designer is to minimise structural weight and keep the external shape as clean as possible in order to minimise drag – the force that seeks to hold the aircraft back as its engines work to push it through the air. This is no academic exercise – every extra pound of weight and pound per square foot of drag represent a penalty for the aircraft operator. In a vicious downward spiral, the engines have to work harder, burn more fuel and emit more greenhouse gas. The airline's fuel bill is higher, ultimately affecting ticket price, passenger numbers and the bottom line.

The problem facing modern businesses is similar – the mountainous weight of their own information is beginning to slow them down, add to their costs and bloat their carbon footprints, with storage provision now set to overtake the processing functions in terms of total energy consumption.

In this sea of data many companies thirst for knowledge. They have no way of properly measuring how much of their storage is being used effectively for valuable business data, how much is hosting data that is no longer used or needed, or even how much of it is free space.

Then there's duplication – how many of the teeming terabytes and petabytes represent multiple copies being held for no good operational reason? Such data inflation is most usually driven by a reasonable fear of losing a unique copy. But the response is unreasonable – copying once, twice or several times. The problem is exacerbated when the same thing happens with development and test data, which proliferates during the life of the programme, when it is genuinely needed, but is not rationalised afterwards when it has clearly outlived its usefulness.

Conversely, even when a company sets out to store offline data that is no longer needed in the day-to-day business, there is no guarantee that this will be done efficiently. The number of copies being held can run out of control, while it can prove difficult or even impossible to retrieve information that later proves to be required.

Highlights

Runaway proliferation of backups is the price to be paid for the benefits of mobile computing

Very often it is the desktop that poses the greatest challenge to an organisation struggling to damp down the data eruption. The operational flexibility yielded by mobile computing is highly desirable, but there's a price to pay. Many companies struggle to manage their estates of mobile devices, resulting in runaway proliferation of backups. In one company, growth in storage of data generated by laptop applications, for example, was recently put at 80-90 per cent a year, and the indications are that 50 per cent isn't unusual across commerce as a whole.

Paying the fuel bill

Like an airline operating a badly designed aircraft, a company that cannot control its data ends up working harder and spending more to achieve less.

Companies pay more than they need for storage, then lose again when it is inadequately managed

For each dollar it spends on fuel the airline generates fewer available passenger-kilometres – the basic unit of air transport productivity – than it would with a more efficient aircraft. The company is similarly penalised, paying more than it needs to for storage and suffering productivity losses because it takes longer than necessary to retrieve information from incoherent, inadequately managed facilities.

There are a number of reasons for this. One arises from that old bugbear "silo thinking," so that a department will acquire capacity for the purposes of a single project, without regard for the wider needs of the company. Very often this is not helped by rigid ring-fencing of departmental budgets and a lack of incentives for managers to create and use shared resources.

Another is an inability, because of inertia in the budget process, to act quickly when extra provision turns out to be needed. This leads IT teams to buy more than they actually require in order to avoid the pain of delays in the future. And then there is the readiness of storage salesmen paid by the terabyte to offer bargain-basement deals on raw capacity.

Highlights

Company accounting practices often offer no incentive to generate and store data efficiently

Possibly the biggest culprit of all is the way that many companies account for IT infrastructure spending, allocating the cost across the organisation without regard for how efficiently or otherwise individuals and teams generate and store data. This itself is often the consequence of the comprehension gap that still yawns between IT specialists and the business functions they serve.

In many cases the businesses have no clear understanding of the capabilities they are buying, while the technology-centric IT professionals feel no need to explain what seems obvious to them. Unable to distinguish among the different classes of storage on offer, and in any case usually feeling no direct impact on their budgets, the businesses tend to tick the box for the highest-performing, enterprise-class service when in fact something less capable and less expensive would meet their needs.

Of course, IT managers very often do have a strong grasp of the core aims of the business, but even then they are unable to make a strong case for accurate internal billing for the use of storage. That's because, having failed to deploy effective capacity management tools, they simply don't know who's doing what with storage with any convincing degree of accuracy. This state of affairs can also arise from deficiencies in the governance of individual projects, which are sometimes launched without having in place all the processes needed to maximise the possibility of success.

These factors can come together to result in massive and costly over-provision of storage. But there's a more insidious effect, which, though harder to quantify, could be hurting companies even more than loose spending. The haphazard accumulation of petabytes of unindexed data means that information can't be found without a laborious, often manual, search. Moreover, the problem gets worse the longer the files have been stored.

Highlights

Managements are waking up at last to the damage that inefficient storage could be doing to their businesses

At first sight it's a gloomy picture, but there are some shafts of light. There's growing evidence that managements are waking up to the problem, recognising that it could imperil their businesses, and beginning to think hard about what needs to be done. So how can companies start lightening their data burden? And what can suppliers like IBM do to help?

On a wing and a prayer

The initial reaction has been to seek quick fixes by, for example, trying to persuade staff to be less profligate with e-mail. But altering human behaviour is rarely simple.

Time and again, when e-mail in-box quotas have been imposed, staff have shown great ingenuity in getting round them as they strive to hang on to information that might one day be useful to the business. The outcome is the very reverse of what was desired – instead of reduced pressure there is a completely unforeseen explosion in local storage, sometimes by as much as 200 per cent.

No storage solution can work if it tries to treat all items of information as if they are of equal value

Such sticking-plaster solutions fail because they seek to treat each item of information in the company's orbit as if it was of equal value to all the rest. In fact, some things are far more equal than others and should be stored accordingly – in a well planned and well promoted e-mail archive, with essential business data secured centrally and indexed for rapid retrieval.

The same approach – identifying, labelling and fast-tracking all the information that really matters – will keep companies flying smoothly through the hurricanes of data coming their way.

Highlights

The IT industry is beginning to move towards an integrated approach to solving the data storage problem

A single sky

One of the prime causes of inefficiencies and delays in the skies above Europe is the fact that air traffic is managed by dozens of different national authorities, each responsible for its own comparatively small block of airspace. The countries recognised the problem some time ago and are moving towards an integrated system under the label Single European Sky. The IT industry and the companies it serves are beginning to develop a similarly joined-up approach to the problems posed by data congestion.

IBM is tackling the issue at the strategic level with its Information Lifecycle Management (ILM) methodology, and in the front line with a wide and growing range of specific tools for the management of data storage.

A great aircraft designer of the past coined the humorous maxim "Simplicate and add lightness." This is the essence of ILM. In their struggles to come to terms with the data deluge, many companies have saddled themselves with systems that are both technically complex and laborious to manage. ILM is a recipe for greater simplicity through consolidation of software and hardware, virtualisation and sharing of resources.

Among the central principles of ILM is the need to select the right storage management tools and apply them effectively. Extraordinarily, in the first decade of the 21st century some household-name corporations continue to rely on Access databases or Excel spreadsheets for the management of enormous estates of storage. But with due acknowledgment to their years of honourable service, these standard-setting applications just aren't up to the challenge of managing exponential data growth. Creating static records, they have diminishing relevance to today's dynamic environments – what worked 15 years ago for small estates is no longer equal to the petabyte challenge.

Fortunately, ILM is supported squarely by a growing array of here-and-now tools for use by companies determined to get to grips with the problem of runaway data.

Highlights

The first step towards a solution is a clear understanding of the nature and scale of a company's data difficulties

The benefits of Process Excellence include optimally loaded storage, enterprisewide good practice, and easier maintenance and training

Pulling out of the spin

Like a pilot out of control in a spin, the company in data difficulties needs above all to understand fully the scale and nature of its predicament. IBM calls this stage of the recovery "Transformation Strategy". Performing the analysis is the job of a suite of offerings from Novus, a company recently acquired by IBM as part of a systematic effort to round out its range of storage management products and services.

The tools do three things. First, they scan the customer's storage environment to produce a "management complexity factor" (MCF) – a measure of how labour-intensive and costly it is to run. Next, they produce a familiar cost-benefit analysis. Finally, they map out potential solutions.

These can be founded upon, among other things, IBM's Enterprise Standardisation Program (ESP) – a set of processes and policies designed to lead to a standardised way of doing things throughout the organisation. There are several benefits to what IBM labels "Process Excellence". Storage could be loaded to its maximum, potentially eliminating wasteful voids. Good practice could be enforced across a variety of different technical environments. Maintenance could prove to be easier, and staff training is potentially less of an effort.

Once a data repository has been brought to order, it needs to be kept that way by developing "Global Insight". That's the task of Storage Enterprise Resource Planner (SERP) – essentially a framework designed to use the available resource management, monitoring and reporting tools to make a better job of running the infrastructure. SERP is designed to push up utilisation rates and remove orphan data, all in a highly automated way and largely from a central point.

SERP is designed to replace today's deeply inadequate practices – with their dependence on spreadsheets and human labour, and very long response times – with a regime that is designed to yield overnight reports that give a consolidated, accurate view of current usage patterns along with insights into future demand patterns. Sometimes described as a "single pane of glass," SERP is effectively a single powerful instrument with which to view all the company's data storage.

Highlights

It is possible to give storage administrators what they need to see exactly what is happening across the network and to take full control of complex environments

Underlying SERP is IBM's TotalStorage® Productivity Centre (TPC) suite of software for the storage layer. This also feeds the initial input into SERP, sitting in a layer below it and providing customised decision-support reports.

Once the storage administrator can see clearly what is happening across a sometimes very disparate network, he needs to be able to take control. Until very recently this had to be done MS-DOS-style, typing in arcane command lines that showed up on a plain green screen. Meanwhile, the world at large had moved on to the clear, intuitive graphic user interfaces that the PC user in the street takes completely for granted.

An IBM product called SAN (Storage Area Network) Volume Controller (SVC) is designed to offer the clarity, ease of use and protection against data corruption and loss that are dictated by ever more complex storage environments. At the core of SVC is virtualisation – creating 'software machines' in order to achieve high levels of flexibility and versatility within the limits of the available hardware. With SVC in place, providing an "abstraction layer," it is possible to overcome the limitations imposed by a rigid correspondence between physical server and physical discs and manage storage in a much more nimble and responsive way.

There is also much to be achieved by improving the storage medium itself. IBM recently acquired the Israeli company XIV, creator of a very advanced storage device called Nextra. This has many of the capabilities described above already built in: it takes virtualisation in its stride and comes with a full suite of state-of-the-art interface tools.

Nextra is also highly scaleable – capacity can readily be added as demand increases – and quick-reacting. Its speed of response is achieved by using a very clever algorithm to spread a given data file in packets over hundreds of highly efficient drives and then retrieving them all simultaneously when the information is needed.

Highlights

IBM's ability to respond to the storage challenge continues to grow

Other recent acquisitions have added to IBM's ability to respond to the storage challenge. They include Diligent (deduplication software), Arsenal (managed backup services for server and PC data) and Softek (migration and portability software to reduce risks when moving data as a result of technology upgrades or business mergers and acquisitions).

Not all of the weapons in the struggle to avoid drowning by data are exclusive to IBM, however. Solutions like Wiki and Sharepoint offer ways of avoiding the local storage of large numbers of copies of the same file. They support and encourage the centralised, Web-style publishing of information so that users have no need or desire to hold their own copies, and in many respects they replicate the central repositories of mainframe days.

Also harking back to the mainframes is the concept of "virtual infrastructure". This can include virtual desktops – essentially the dumb terminals of old, but with full, user-friendly graphics.

Petabytes of storage can now be managed as a single entity

Supporting the concept of Web-style publishing of information, IBM's Scale-Out File Services is designed to provide a single name-space for Web content management. It runs on open-source operating systems like Linux®, using IBM's highly versatile and scaleable General Purpose File System, and is designed to allow petabytes of storage to be managed as a single entity rather than as a multitude of scattered "islands".

Highlights

Companies buying data storage need a classification system to help them spend their money appropriately

Right class, right price

For all their many points of resemblance, there is still one big difference between air transport and IT. While consumers of air travel completely understand the different classes of service and the ticket prices attached to them, there is very little such differentiation for the business user buying data storage. There is so little financial penalty attached to adopting a full, enterprise-scale solution that this is done as matter of course. As a result, the company often pays more than it needs, and its IT team finds itself managing a Boeing 747 level of complexity when the business need would have been met by a ten-seat turboprop.

Until recently, businesses would not have been able to buy lower and more appropriate levels of service. But now IBM is promoting a model similar to that of the airlines – no more than half a dozen classes of data storage service, each clearly defined in terms of price, levels of back-up and security, speed of response, and provision for both routine and disaster recovery. The more critical the data to the affairs of the company, the higher the class of storage required. The higher the price, too, so that no-one in the company is in any doubt that a good case needs to be made for buying higher-specification provision in preference to something simpler and cheaper.

Simplicity is essential to the success of such an approach, and that means developing straightforward, readily understood criteria for the classification of data into the various classes – from "economy" for routine housekeeping information to "super first-class" for things like the company's core intellectual property. If the process of labelling data is complex, or if there are too many classes, a bad situation could become even worse.

What's needed to make data classification and rational, appropriately priced storage regimes a practical reality? IBM believes there must be a formal, industrywide effort in co-operation with leading users to develop a standard classification capable of winning broad support.

Highlights

Chief information officers must be empowered to drive through data classification, users must pay according to the level of service they receive

It's one thing to publish a standard and quite another to see it take effect, however. Translating worthy intent into real behaviour will call for a number of measures within companies. These could include expanding the role of the traditional chief information officer, fully empowering him to drive through measures like data classification, and changes in internal accounting practices to recognise the different levels of service and ensure that users insisting on first-class pay significantly more than those who are content with economy.

If multi-class data storage becomes a mainstream practice, the users can expect to see their costs per unit of storage reduced and their information retrieval times improved. But there should be more good news. A more efficient storage environment is also cheaper in terms of its impact on the environment, drawing less power and thus demanding less of CO_2-producing generating systems.

Energy-efficient hardware and better management of resources will help to cut environmental damage

IBM has developed an approach called SPACE to help users understand the environmental impact of their existing storage infrastructure. Replacing it with more energy-efficient hardware – just as the airlines are now racing to phase out their old kerosene-guzzling aircraft and introduce designs offering a 20 per cent improvement in fuel consumption – will yield an immediate payoff. That in turn will be magnified by the fuller and more efficient use of resources that will result from applying the tools and disciplines described earlier.

Highlights

The data storage community can emulate air transport and develop a cleaner, more efficient way of operating

Clearer skies ahead?

Air transport and the data storage sector of IT are beleaguered industries. Both are under the cudgel to undo the consequences of excess – too many aircraft movements, too many clogging terabytes of ill-organised information, too much delay, too much cost, too many tons of climate-changing CO_2.

Aviation is rapidly gathering speed in pursuit of a cleaner, far more efficient way of operating – and so too can the data storage community. Companies like IBM know what needs to be done, have many of the means to do it and are working diligently to develop more. All that's missing now is an industrywide determination, to be followed by systematic co-operation, to put in place a set of standards that could help users to acquire storage prudently, rein in runaway data growth and manage their precious information more effectively for the greater good of their enterprises.

Highlights

Recent IBM successes

IBM wins

Using SERP to identify reclaimable storage has allowed one major US company to cut spending by many millions of dollars. Other companies have each saved tens of millions of dollars.

A Fortune 50 company used SERP to identify enough orphaned storage to pay for the product and its implementation twice over.

A global bank used IBM-Novus products to save tens of millions of dollars a year on Tier 1 storage. A Fortune 500 company identified reclaimable storage worth $7 million.

An insurance company applied a suite of IBM storage products and services and achieved a 500 per cent return on initial programme cost.

A US hospital found that its storage costs fell by half and storage capacity grew 500 per cent after the installation of an IBM disk system.

Cost of ownership halved as a result of increased utilisation when a Fortune 500 power company implemented an IBM tiered-storage solution.

A US law enforcement agency deployed an IBM-based nationwide identification application and cut the cost per suspect identified by more than 80 per cent.

A US regional bank was running into trouble with data storage – its stock was doubling every 12-18 months and disk utilisation was a very uneconomic 28 per cent. The addition of storage visualisation software from IBM boosted utilisation to 80 per cent and slashed costs.

A leading state healthcare insurer used IBM technology to consolidate its storage, improving performance threefold and cutting backup times by 50 per cent.

A global food company put 15 months of data, equal to three million documents, into an IBM-based archive. Back-office processes ran 20-25 per cent faster and there was an initial cost saving of $70,000, with more in prospect.

More information

For more information on infrastructure solutions from IBM Global Technology Services, please visit:
ibm.com/services/gts/uk
or e-mail:
chris_mackey@uk.ibm.com
tonyr_cox@uk.ibm.com

IBM United Kingdom Limited
76 Upper Ground
South Bank
London
SE1 9PZ

Contributors

Chris Mackey – Chris is an Infrastructure Consultant and has worked across a broad range of industries in service delivery, infrastructure and processes.
E-mail: chris_mackey@uk.ibm.com

Tony Cox – Tony has extensive infrastructure services experience across both systems and storage management in both Specialist Sales and Consulting roles. He helps clients address their information and storage management challenges.
E-mail: tonyr_cox@uk.ibm.com

Simon Richardson – Simon is an IBM Commercial Manager and has extensive client-facing experience, including work as an Infrastructure Storage Consultant.
E-mail: Simon_Richardson@uk.ibm.com

SDW03010-GBEN-00