

Cost of Losing Information

A Framework for Information Management Planning

January
2005



1601 Trapelo Road
Waltham, MA 02451
1.866.WANSync
www.xosoft.com



Executive Summary

What is the proper concern of the IT department of a modern enterprise – backing data up or making sure that data is available? The answer to that is obvious – the primary information management concern in the enterprise today is to ensure that the knowledge necessary to drive critical business processes is available where it needs to be, when it needs to be.

The costs of failure to do this are high. A recent study of 80 large organizations by Infonetics Research found that overall downtime costs averaged an astounding 3.6% of annual revenue! In another study, Forrester estimated the *average* cost of downtime for e-commerce sites at \$8,000 per hour – at larger sites like eBay, Intel and Amazon, the costs soar to hundreds of thousands of dollars per hour. And research by Creative Networks, Inc. has shown that Microsoft Exchange downtime can also easily cost thousands of dollars per hour in a mid-sized organization.

On another end of the spectrum are the potential costs of legal exposure arising from losses of data covered by legislation like HIPAA or Sarbanes-Oxley. Speed is less an issue here, but the cost of losing the information entirely could be devastating.

Real-time information management is about getting information where it needs to be, when it needs to be, whether that is microseconds, minutes, or days.

Business in the real world requires that information be accessible in real time – in fact, Gartner has recently begun to emphasize the importance of a real-time infrastructure. But what does real time really mean?

Real-time is usually associated with speed – a real-time system must be able to process information quickly and as events actually occur in order to respond in a timely way. Real-time doesn't actually mean fast – it means *fast enough*. What is fast enough for real-time response to one event, say a pilot's routine request for the status of some subsystem, is woefully inadequate for another, for example the need to take evasive action to avoid collision with another plane.

In other words, real time information management is about just what we said above: getting information where it needs to be when it needs to be, whether that is microseconds, minutes, or days.



From Data to Information to Knowledge

Finding adequate information management solutions requires understanding just what it is we are managing. Let's start from the bottom – the raw data.

Data is nothing more and nothing less than a collection of bits stored on some medium, a tape, disk, or CD. A sequence of random numbers is data, as is the collection of bits representing your current inventory. In and of itself, data has no real meaning, and hence no real value.

Until we understand how it is organized.

Information is organized data – and the organization gives the data meaning. When a collection of bits is understood to be ASCII text or a digital image, it gains meaning and, therefore, value. In the information technology world, this organization comes from applications –applications impose structure on data and thereby create meaning, and value.

Information that sits isolated is not especially useful. The tree may make a sound when it falls in the forest alone, but it has no bearing on anyone's life. Similarly, information only has impact when it is connected – connected to context, to other information, and most importantly, to people. This is the next level.

Knowledge is connected information and it is knowledge, therefore, that can actually cause things to happen, bringing the highest level of value.

In the end, only knowledge matters, because only knowledge can have any impact on your business. And so the task of information management is essentially to keep your information connected so that it becomes a part of the knowledge that drives your business. This is just a different way to say what we said at the start: the value of information for your business, and therefore the driver of your information management efforts, is ensuring that the right information is kept connected – or available – to those who need it, when they need it.

The Value of Knowledge

An obvious approach, then, is to maintain total connectivity – to keep all information online and immediately available to all users all the time.

This can be a fairly daunting task even in a static environment. In today's highly dynamic environment, with new information generated at explosive rates, a total connectivity approach is impractical – the cost of the required hardware, software and administrative resources is prohibitive.

What is needed is to align IT resources with the actual value of the information being supported. How can this be done?

There are two different kinds of value that can be distinguished for information. We are most familiar with *intrinsic value* – the value of the information itself. While an absolute scale for measuring intrinsic value is problematic, it is straightforward to obtain a useful quantitative measure simply by considering what it might cost, in the worst reasonably likely case, to completely lose the information. This value may, of course, vary with time.

Information only has impact when it is connected – to context, other information, and especially to people.



A different, but equally critical kind of value is the *urgency* value, which depends on the time required to access the information. It can be measured practically by estimating the cost of a delay in accessing the information.

Two examples may help clarify the difference between intrinsic and urgency value.

The database behind a heavily-trafficked e-commerce application or a large Exchange server has a very high urgency value – a delay of an hour can easily cost tens of thousands of dollars. The intrinsic value of the database, particularly of individual transactions in it, is also high, but decays over time, dropping when the daily backup is done, and dropping again precipitously after the transaction has been processed through the financial and shipping systems or has been passed on to its intended recipients.

In contrast, historical financial records covered by the Sarbanes-Oxley act for a publicly traded company have a high intrinsic value for as long as they are required to be stored, since the company is exposed to a significant legal risk if they are lost – after this time has expired, of course, the intrinsic value falls to nothing. Their urgency value, on the other hand, remains relatively low throughout.

Real Time, Real World: Information Lifecycle Management

A real-world, real-time system must manage to both intrinsic and urgency value. The intrinsic value determines the required level of certainty that data will be protected and available – this is typically accomplished through various kinds of redundancy. Managing to urgency, however, means allocating IT resources so as to support rapid access to urgent information, and low-cost access to less urgent information. To repeat the theme, an effective information management system gets information where it needs to be when it needs to be there – not before, and certainly not after.

The Information Lifecycle

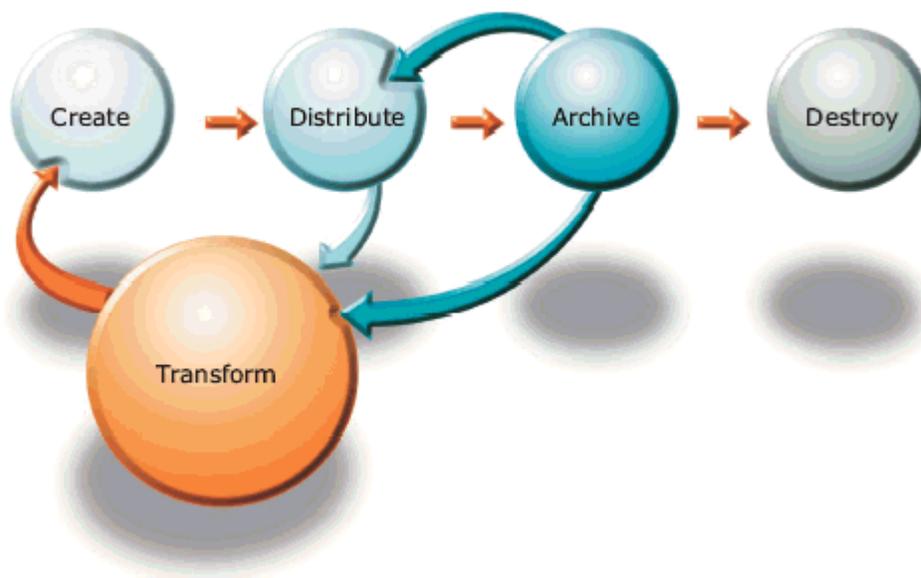


Figure 1: Information Lifecycle Management - managing to urgency



One implication is that the resources required to connect a given item of information can change over time as its intrinsic and urgency values change. This realization underlies the concept of *information life-cycle management*. Information has a well-defined lifetime: it is created, has a period of useful life, possibly in several different phases, and eventually loses its value and can be destroyed. The goal is to manage the information appropriately through the course of its existence, avoiding over-investment in protecting and distributing non-critical information, while investing significantly in critical, urgent information.

The information life-cycle management strategy, like any effective IT strategy, must be driven by business-centric policies spanning all applications, processes and resources relevant to a business process, and it must also be managed in such a way that information assets on all platforms and operating systems pertaining to a business process are connected, visible, and manageable as a whole. Most importantly, an effective strategy must align resources with the real value of the data at any given point.

The tasks of information lifecycle management are not fond hopes for the future, they are critical imperatives today.

Eventually, information life-cycle management may be significantly simplified and fully automated, however, it is important to understand that the tasks of information life-cycle management are not just fond hopes for the future – they are critical imperatives.

So let's turn to some practical considerations for dealing with these tasks.

A Framework for Information Management Planning

The fundamental purpose of business is to create value; it also has a responsibility to protect the value it creates. These two imperatives should drive the information management system.

In fact, value is the element that brings together all that we have been discussing into a general framework for evaluating the components of an information management solution.

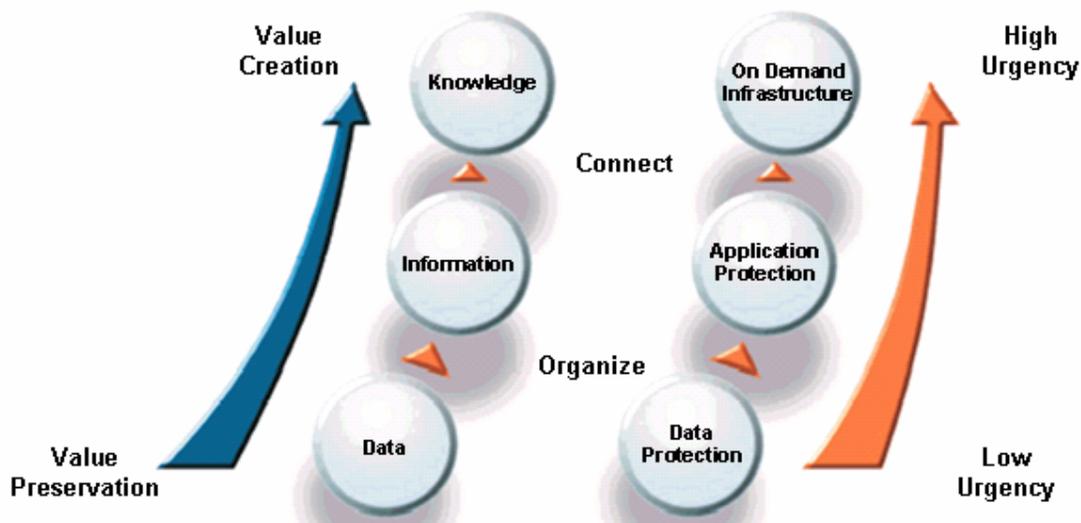


Figure 2: Business activities in terms of their relationship to value



In the preceding figure, there are four parallel segments, each increasing from bottom to top in value and complexity. The left-most segment suggests the spectrum of business activities in terms of their relationship to value, whether they preserve value that already exists or creating new value. Balance is important here, but it is clear that a business cannot be successful without a distinct bias toward the creation of new value. The right-most segment echoes this – information that is urgent is more likely to be associated with activities that create value.

This is clearly reflected as well in the second segment, which illustrates the hierarchy of data, information, and knowledge introduced earlier. Storing data may be a critical task in protecting value that a business has already created, but the ability of knowledge to impact things is key to creating new value.

The new element in the figure is the third segment, which frames information management technology in terms of business value.

At the bottom are traditional data protection technologies. These include tape backup, offsite storage, and archiving solutions. More recent technologies, like snapshot and disk backup, also fall into this category since they operate within the same model. These technologies are largely value-preserving and are typically appropriate only for less urgent information.

The middle level represents protection technologies that center on applications rather than data. These obviously include capabilities built directly into applications, for example, the ability to detect and repair errors resulting from unexpected shutdown. Versioning at various levels of sophistication falls into this category, as do application-aware migration and replication technologies, and offline-processing solutions. These solutions typically support faster access and thus support more urgent information – they can be either value-preserving or value-creating.

Are you investing in systems that protect information in a way and at a cost that are aligned with its value?

The top category has been described in a variety of ways, from ubiquitous computing to on-demand computing. It consists of technologies that go beyond simple protection of data or applications to actively enhance business value, generally by adding connectivity. Examples include high-availability solutions, which maintain access even through application or hardware failure; content delivery solutions, which proactively deliver information where it is needed; and collaboration solutions, which go beyond connecting a user to information to connecting users to each other as well. Obviously, these solutions typically deal with information that is very high on the urgency-value scale, in some cases providing access even as information is being created.

This framework should guide your information management infrastructure planning. Are you investing in systems that protect information in a way and at a cost that are aligned with its value? Are you leveraging resources not just to protect value, but to help your business create it? And, at the detailed level, what technology is appropriate to a given information management task – simple data protection, application protection, or truly adding business value? These considerations provide an invaluable guide to allocating scarce IT resources.



About the Author

Eric Jackson is Vice President of Products at XOsoft. Mr. Jackson has nearly 20 years experience in the development and commercialization of advanced software technology.

In 1986, he joined the staff of the Applied Math program at Princeton University. His work there focused on the design of software to solve complex turbulent flow problems on advanced supercomputer and parallel supercomputer architectures and research into turbulent flow phenomena. Mr. Jackson has numerous publications in scientific journals, including *Nature*, *Physical Review*, *Physical Review Letters*, *Proceedings of the Royal Society of London*, *Proceedings of the SPIE*.

In 1994, he moved to a private consulting firm in Princeton. There, collaborating closely with researchers at leading computer chip manufacturers such as LSI Logic, Toshiba, DEC, IBM and others, he designed and implemented one of the first and most sophisticated software applications to automatically analyze and correct computer chip designs to improve their manufacturability.

In 2000, he co-founded Ibrix, Inc. He is the inventor of the Ibrix distributed file system, a parallel file storage system able to scale in size and performance to millions of terabytes. Mr. Jackson served as chief software architect and led the product development effort through the initial phase of the company's existence.

In 2003, Mr. Jackson founded DeepWeave, a management consulting practice helping leading-edge software companies with technology commercialization and strategic product development.

About XOsoft

XOsoft's continuous application availability solutions keep enterprises working. The company's patented business continuity, disaster recovery and continuous data protection software platforms significantly lower the risks of decreased productivity and lost revenue by minimizing both planned and unplanned critical application downtime. Many of the world's largest corporations and government agencies have chosen XOsoft to ensure uninterrupted access to their most valuable information assets. For more information, please visit www.XOsoft.com.

A 14 day trial evaluation of the software is available for download at:
<http://www.xosoft.com/download/index.shtml>