

RECOVERABILITY AND REGULATION IN THE U.K.

BUSINESS VALUE WHITEPAPER

Double-Take Software, Inc.

Published: June 2007

Abstract

If you are on the consumer side of the thousands of regulations imposed by the British government, just knowing they exist probably gives you some peace of mind. However, if you are on the compliance side of these complex and ever-changing regulations, you know that staying compliant is a constant tug-of-war between resources and regulation. Acting responsibly on behalf of your consumers and ensuring that your company avoids fines brought on by non-compliance means staying on top of regulatory legislation. Using software products for data replication, application availability and system state protection from Double-Take Software can reduce the costs of managing a data protection strategy while increasing confidence that your systems remain compliant and your data and applications remain available.

Introduction

If you are working in a regulated industry in the United Kingdom, you know that there are thousands of regulations concern records management, and thousands more apply if you are sending or receiving electronic data across borders. If you are on the consumer side of these regulations - that is, if you are receiving health care, taking medication, have a bank account, credit card, investment accounts or will rely on any kind of government agency for benefits or protection, knowing these regulations exist probably gives you some peace of mind. However, if you are on the compliance side of these thousands of complex and ever-changing regulations, you probably have already had more than one sleepless night. Acting responsibly on behalf of your consumers and ensuring that your company avoids fines, or worse, brought on by non-compliance, means staying on top of regulatory legislation.

If you are in any part responsible for compliance in your organization, you already know that government regulations don't spell out what technology needs to be in place. For example, regulations may never even mention the word "software"; but, for most regulated companies, necessary audit trails require some kind of record management software. Technology changes so quickly that specific requirements can't be legislated. In addition, legislators also assume that an organization will comprehend the requirements and implement a right-sized solution for their particular situation, service or product.

Compliance and Disaster Recovery

Just as government regulations require an audit trail without defining the technology to be used, they all require a disaster recovery (DR) plan for certain kinds of data. Again, legislators urge steering members of the organization to read regulations for comprehension and then implement a solution that best reduces risk, protects privacy, ensures accountability, and in some cases, ensures business continuity in a disaster.

As business in the UK becomes more of a global industry, institutions are constantly reviewing overseas policy – as it may become part of local legislation and best practice policy. Below are major regulations from the UK and the US that require companies to safeguard information - and produce it on demand in an audit.

Major U.K. Regulations

Retail Mediation Activities Return (RMAR)

The RMAR requires anyone who has responsibility for collecting and submitting data to the FSA (Financial Services Authority) to produce management accounts and input the details onto the RMAR section of firms within thirty days. This is required twice a year and companies who fail to meet this deadline will incur fines. In some cases, as was the case for 20 companies in 2006, a business can simply be stopped from trading if they fail to comply or complete the RMAR.

Information Assurance (IA)

Information Assurance is a policy and guidance framework for electronic government progresses, including development, procurement, provision and maintenance of governmental (both central and local) services. IA drives to implement policies and tools to enable and assure the availability, integrity and confidentiality of e-Government services, having particular focus upon identity registration, enrollment and authentication processes, which underpin access to these services.

National Privacy Principle 4

This principle, part of the Guidelines on Privacy in the Private Healthcare Sector (and NHS), requires that health care providers take reasonable steps to protect the health information it holds from mis-use, loss, unauthorised modification or disclosure and the ability to destroy data. Personal information is required to be managed, kept current and secured so that the threat of identity theft is reduced.

National Privacy Principle 6

Under this principle, a health service provider must have an environment that is documented, and policies to control it, so on request a service provider can provide the information requested on health management or personal information that is stored. This is a legislative requirement and the timescales that are permitted for a health service provider to respond are extremely short.

Data Protection Act (DPA)

The Data Protection Act provides privacy and protection of personal data for consumers in the U.K. The act states that "appropriate technical and organisational measures shall be taken against unauthorised or unlawful processing of personal data and against accidental loss or destruction of, or damage to, personal data."

Implications of U.S. Policy on U.K. Regulation

UK healthcare, government and financial services providers are beginning to adopt many rules and regulations from the U.S. Though these regulations are currently without the stiff penalties imposed in the U.S., they are being adopted for best practice. For example, National Privacy Principle guidance relies on the U.S. Health Insurance Portability and Accountability Act (HIPAA) for an extra level of consumer protection. Additionally, corporations with more than 300 U.S. shareholders must comply with Sarbanes-Oxley regulations.

Major U.S. Regulations

Sarbanes-Oxley Act

In 2002 the United States federal government passed the Sarbanes-Oxley Act (SOX) which establishes laws and standards for U.S. public company boards, management, and accounting firms. Provisions include a requirement that public companies evaluate, disclose and qualify (by independent auditor) the effectiveness of internal controls for financial reporting, a ban on personal loans to any executive or director, prohibition on insider trading during certain periods and accelerated reporting of insider trading, protections for whistle blowers and increased penalties for security violations. Because financial reporting in most companies is supported by electronic systems, IT is a large part of internal control.

SOX and DR Planning

Section 404, Management Assessment of Internal Controls pertains to disaster recovery requirements, stipulating that an organization should:

- State the responsibility of management for establishing and maintaining an adequate internal control structure.
- Contain an assessment, as of the end of the most recent fiscal year of the issuer, of the effectiveness of the internal control structure and procedures of the issuer for financial reporting.

Translated, DR planning for SOX has two primary parts: implementing systems that completely protect all financial and other data required for reporting regulations and providing data on-demand, and clearly documenting those procedures so auditors can readily see that the plan protects regulated data as required.

FDA Title 21 Part 11

The U.S. Food and Drug Administration (FDA) is a public health agency that protects American consumers by enforcing laws that regulate the manufacture, storage, import and sale of food and medicine for humans and animals, as well as medical devices and cosmetics. CFR Title 21 Part 11 Electronic Records; Electronic Signatures (Part 11) is the rule regarding how all companies regulated by the FDA must maintain electronic records in order to remain compliant with Good Clinical, Laboratory and Manufacturing practices (GxP). Though finance and planning are excluded from Part 11, all other functional areas of FDA-regulated companies that involve GxP must comply or face legal sanctions and even criminal charges.

Title 21 Part 11 and DR Planning

Part 11 has two basic requirements: that companies are able to generate accurate and complete copies of records for inspection and review during an audit, and that those records are protected so that they are readily retrievable throughout the required retention period. Data availability and protection at this level require a good business continuity plan that takes into consideration high availability, or "failover", and disaster recovery. The FDA realizes the potential impact of interruptions to critical business applications and therefore requires that electronic GxP applications must be always available – or at least quickly recoverable. For example, if a Windows-based application involved in the supply chain for a critical drug fails - interrupting the national supply of the drug - millions of lives are at risk. Another goal of continuous business operation is disaster recovery – the ability to restore critical data and operations, to new hardware at the original site or at a different physical site, after a disaster.

Gramm-Leach-Bliley Act

The Gramm-Leach-Bliley Act (GLBA) governs consolidation of financial institutions and implements financial privacy rules and safeguards. The GLBA governs how customers' personal information is collected and disclosed and requires safeguards to protect this information. The regulations also apply to any company that receives this information, even if they are not financial institutions.

GLBA and DR Planning

The Safeguards Rule contained in the Gramm-Leach-Bliley Act requires regulated institutions to:

- Insure the security and confidentiality of customer records and information.
- Protect against any anticipated threats or hazards to the security or integrity of such records.
- To protect against unauthorized access to or use of such records or information which could result in substantial harm or inconvenience to any customer.

Health Insurance Portability and Accountability Act

The Health Insurance Portability and Accountability Act (HIPAA) make provisions for employees and their families to keep their health insurance when they change or lose their jobs. HIPAA also established national standards for security and privacy of electronic transmission of health care data.

HIPAA and DR Planning

HIPAA's security rule states that each organization must determine its own risk in the event of an emergency (that would result in loss of operations). However, there are three things that companies must be able to substantiate in regards to DR planning:

- There has been a formal analysis of risk to data (physical and virtual access).
- A DR plan exists that covers backup, storage and recovery.
- The DR plan adequately addresses the risks outline in the analysis.

Data Protection and Availability

Data protection and data availability are the two major goals of a backup, availability and disaster recovery plan. The quantitative measures used are called Recovery Time Object (RTO) and Recovery Point Objective (RPO).

Recovery Time Objective (RTO) and Recovery Point Objective (RPO)

RTO represents the amount of time between the start of an outage and the resumption of normal business operations. The RTO for an outage that can be resolved by reloading from tape backup includes the time necessary to locate and mount the tape, the time required to restore the data from the tape, and any time necessary to post-process the restored data before restarting the downed applications.

Recovery Point Objective (RPO) represents the point to which business data can be restored. This can be thought of as the latency between a production data set and its redundant or replicated copy. This latency may be expressed as a number of changes or a time interval; it measures how out-of-date the replicated copy will be compared to the original. For example, a nightly backup means that the RPO will be the time between when data was written to tape and when the failure occurs: a failure any time on Tuesday has an RPO of Monday night.

Acceptable RTO and RPO must be evaluated by each enterprise individually. However, an RTO of less than 24 hours is usually necessary in order to ensure that supply of critical products or services to consumers is not interrupted. RPO can much more time-sensitive when it comes to compliance. If an application is saving time-stamped electronic signatures, there can be no gap in protection. While tape backup systems can be inexpensive and fairly reliable, the RPO and RTO they offer is not enough to remain compliant.

Loss Scenarios

Successful IT organizations evaluate their disaster recovery and business continuity options using these measures in the context of the types of loss for which they need to plan. There are four typical scenarios which should be evaluated when putting together a disaster recovery plan and selecting software or hardware solutions as part of that plan. Just as each system should have its own RPO and RTO goals based its criticality to the organization, each system should also be evaluated in terms of potential loss scenarios. By understanding RPO and RTO in the context of potential loss scenarios, an IT organization can more effectively plan for recovery of the system.

- **Loss of a single resource** – a single server supporting a regulated functional area fails or is interrupted. For example, losing the server that supports the document repository would seriously affect the ability to provide current documentation in an audit. It would also affect the ability to maintain the time-stamped electronic signature for document versions.
- **Loss of an entire facility** – an entire facility, and all of its resources, is unavailable due to natural disaster, power outages, failure of the facility's environmental conditioning systems, or terrorist action. Unless the facility in question is the primary production facility, the best response is to resume normal operations at another site.
- **Loss of user data files** – accidental or intentional loss of critical data files. The best mitigation is to restore the lost data from backup - normally, from the previous RPO. If there is a gap between backups, time-stamped data that is necessary to establish an audit trail can be lost.
- **Planned outages for maintenance or migration** – the need to restore or repair service. If this activity is not transparent to users and requires forced downtime, productivity will slow or stop during the outage.

Providing Effective Disaster Recovery

Depending on the crisis that drives the recovery, DR may take several different forms. In the most complex scenario the complete failure, destruction, or interruption of access to a data center might necessitate moving the company's operations and personnel to an alternate set of servers at another location. More simple recoveries might involve restoring operations after damage to the primary copy of critical data.

There are a number of strategies that can be employed to protect important data, and each has strengths and weaknesses. The most common method of storage protection is also the oldest: backing up to and restoring from magnetic tape. This method has been around for almost forty years and is still the bedrock of most recovery strategies. The cost per megabyte for tape storage is low; it's easy to move tapes to secure offsite storage, and the technology continues to scale well for many applications. However, tape backups have limitations, such as the amount of time required to back up and restore large volumes of data, the accompanying latency between when the data was protected and when the loss occurs, and the security involved in moving tapes to offsite storage. Accordingly, much attention is being focused on replication-based technologies.

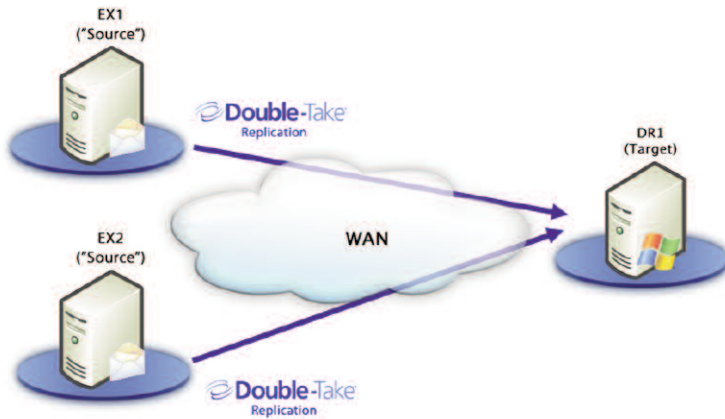
Replication-based technologies offer the promise of capturing a data set at a particular point in time with minimal overhead required to capture the data or to restore it later. There are four main methods of interest in today's storage environments:

- **Whole-file replication** copies files in their entirety. This is normally done as part of a scheduled or batch process since files copied while their owning applications are open will not be copied properly. The most prevalent use of this technology is for login scripts or other files that don't change frequently.
- **Application replication** copies a specific application's data. The implementation method (and general usefulness) of this method varies dramatically based on the feature set of the application, the demands of the application and the way in which replication is implemented. This model is almost exclusively implemented for database-type applications.
- **Hardware replication** copies data from one logical volume to another and copying is typically done by the storage unit controller. Normally, replication occurs when data is written to the original volume. The controller writes the same data to the original volume and the replication target at the same time. This replication is usually synchronous, meaning that the I/O operation isn't considered complete until the data has been written to all destination volumes. Hardware replication is most often performed between storage devices attached to a single storage controller, making it poorly suited to replicating data over long distances. Most hardware replication is built out of SAN-type storage or proprietary NAS filers.
- **Software replication** integrates with the Windows® operating system to copy data by capturing file changes as they pass to the file system. The copied changes are queued and sent to a second server while the original file operation is processed normally without impact to application performance. Protected volumes may be on the same server, separate servers on a LAN, connected via storage-area network (SAN), or across a wide-area network. As long as the network infrastructure being used can accommodate the rate of data change, there is no restriction on the distance between source and target. The result is cost-effective data protection.

Supporting Your Compliance Strategy with Double-Take

Double-Take from Double-Take® Software can fulfill the backup, disaster recovery and emergency mode operations that are required for compliance. Double-Take is a real-time data replication and failover application that augments an existing network environment by providing a data protection mechanism that has minimal impact on users or network resources.

Double-Take allows the administrator to specify that mission-critical data stored on a network server should be protected by creating a second copy of the data on another system, usually at a disaster recovery site. Double-Take monitors any changes to the production copy of the data and replicates those changes to the secondary server. This second copy of the data is synchronized in real-time with the first, making the data accessible in the event of a major disaster or system outage.

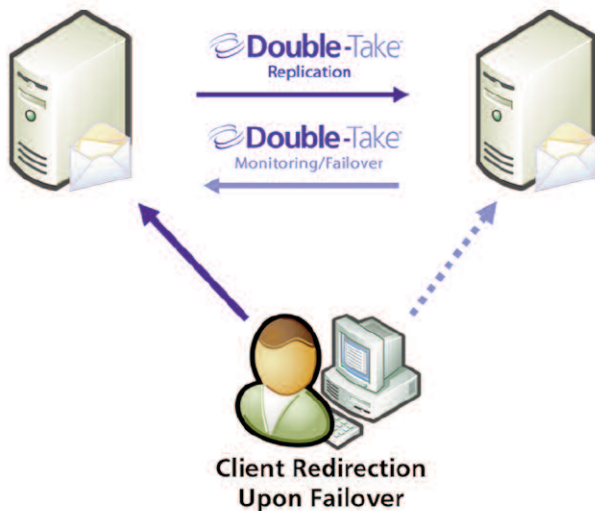


Disaster Recovery

conditions requires us to examine three scenarios in which replication might lead to better business continuity and disaster recovery in support of an organization's compliance efforts: local/remote availability, disaster recovery and backup and restore.

Disaster recovery is often assumed to be synonymous with business continuity. However, true disaster recovery is the ability to restore needed business data after a disaster. Many administrators and planners think of disaster recovery as the ability to quickly resume operations at a separate physical site; that's an overly broad and very expensive capability that's mostly relegated to very large organizations whose businesses are able to justify the cost and logistical complexity. Double-Take gives you large-enterprise disaster recovery abilities at a fraction of the cost of large, dedicated "hot site" recovery operations. The ease of administration and high scalability of Windows and Double-Take means that disaster recovery is more affordable, easier to deploy, and simpler to operate than other DR solutions.

One of the most common approaches to continuous business operations is that of "failover". The goal for these high availability (HA) solutions is to keep the users productive even when outages affect their servers. HA is often thought of in terms of implementations of highly redundant hardware and Windows clustering; however, in many environments file and print servers are perceived as not worth the expense and complexity of a true HA solution. Double-Take's ability to copy needed data from multiple sources to a single target provides a scalable and cost-effective way to enhance the resilience file servers, for example. Three key benefits from leveraging Double-Take protection for protecting critical applications on Windows servers are:



Local/Remote Availability

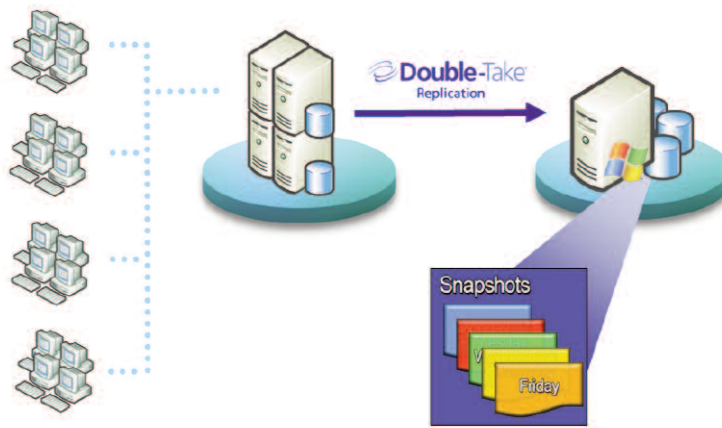
and applied to the target data sets.

- A single Double-Take target can handle failover for multiple source servers simultaneously. If a source server fails, its workload can be

The greatest strength of the replication technology leveraged by Double-Take is that it can operate efficiently by only replicating the data that's changed. Combined with bandwidth throttling and queuing, this allows software replication to work well over long distances, even with slower WAN links.

The approach of replicating data in real time offers a potential escape from the cost-versus-recoverability dilemma. The phrase "business continuity" covers a broad spectrum of technologies, processes and planning approaches. Evaluating the usefulness of replication for particular

- Windows servers running Windows NT 4.0, Windows 2000, and Windows 2003 and Windows Storage Server 2003 NAS devices can be equally protected, either to the same target or to multiple geographically distributed targets. This allows quick deployment of a single target for immediate protection, followed with later addition of more targets to improve resilience.
- Windows servers, due to technologies such as Active Directory and NTFS, are easy to manage with familiar tools. They integrate seamlessly into the existing domain infrastructure. User rights and permissions are easily maintained by real-time replication



Enhancing Backup and Restore

redirected transparently to the target; if other sources fail later, their work can also be directed to the target. The high scalability of Windows means that a single device can handle failover of multiple servers without additional administrative overhead.

For a surprising number of companies, tape backup continues to be their only preparation for business continuity. The challenge with this approach is the ever-increasing restore times driven by the growth in data volume and change rates. Consider a typical scenario involving offsite storage and assume that full backups are done every weekend, with nightly incremental backups. Off-site storage is used for continuity protection. A failure that occurs at 4 p.m. Tuesday must be recovered with the previous weekend's full backup and the Monday night incremental -but if that tape has already gone offsite, it must be retrieved which can add hours (if not days) to the recovery time. Even if the tape can be retrieved with only a four-hour lead time, that still means that users won't have access to the Monday version of their data until sometime on Wednesday (and Tuesday's data is completely lost). For many companies, this is not practical. Let's examine data protection strategies and their ability to address the inadequacy of existing tape-based backup solutions:

- Whole-file replication does provide a second copy for backup purposes within the latency parameters discussed earlier. However, most solutions do not properly handle a situation where the target data is being actively backed up. If the backup software locks files as it backs them up, replication may fail until the files are unlocked again.
- Similarly, most application replication tools do not deal well with the target data set being locked for backup.

As with disaster recovery, hardware and software replication offers approaches that are more flexible. Most hardware replication solutions offer various backup enhancements, including freezing one set of data while the other is given over to the backup (which may be host- or storage-attached) and making "snapshot" or point-in-time copies of the data. The only potential caveat is the re-synchronization time required for the frozen data set once it's thawed and updates are allowed to happen.

Software replication via Double-Take can offer similar benefits with a different twist. Unlike hardware solutions where one logical copy of the data exists in two arrays, the two data sets in software replication are only loosely coupled. This means that while the production data is locked and in use, the redundant copies are natively in a closed state (except of course when each file is actually being updated). During the remainder of the time, tape backup software has easy access to the replicated copies, and they can be backed up without placing any additional network or CPU load on the production server; and without the need for expensive backup agents. Changes will continue to be sent from the source to the target and applied after the tape backup is complete.

Summary

Major British and U.S. regulations establish standards for storing, using and maintaining electronic data, but safeguarding information and producing it on demand in an audit is one of the biggest technology challenges that regulated organizations face. Replication-based technologies offer the best strategy for maintaining critical data, with minimal overhead needed capture the data or to restore it later. Double-Take from Double-Take Software can fulfill the backup, disaster recovery and emergency mode operations that are required for compliance. Double-Take is a real-time data replication and failover application that augments an existing network environment by providing a data protection mechanism that has minimal impact on users or network resources. Double-Take can deliver a comprehensive portfolio of services to help assess, design, plan and implement effective data availability and disaster recovery solutions.

For questions on Double-Take, including pricing and product features call toll free 888-674-9495 or send e-mail to info@doubletake.com.

About Double-Take® Software

Headquartered in Southborough, Massachusetts, Double-Take® Software (Nasdaq: DBTK) is a leading provider of affordable software for recoverability, including continuous data replication, application availability and system state protection. Double-Take Software products and services enable customers to protect and recover business-critical data and applications such as Microsoft Exchange, SQL, and SharePoint in both physical and virtual environments. With its unparalleled partner programs, technical support, and professional services, Double-Take Software is the solution of choice for more than ten thousand customers worldwide, from SMEs to the Fortune 500. Information about Double-Take Software's products and services can be found at www.doubletake.com.

© Double-Take Software. All rights reserved. Double-Take, GeoCluster, and NSI are registered trademarks of Double-Take Software, Inc. Balance, Double-Take for Virtual Systems, Double-Take for Virtual Servers and Double-Take ShadowCaster are trademarks of Double-Take Software, Inc. Microsoft, Windows, and the Windows logo are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. All other trademarks are the property of their respective companies.

Double-Take Software Headquarters

257 Turnpike Road
Southborough, MA 01772
Phone: +1-800-964-0185 or +1-508-229-8483
Fax: +1-508-229-0866

Double-Take Software Sales

8470 Allison Pointe Blvd. Suite 300
Indianapolis, IN 46250
Phone: +1-888-674-9495 or +1-317-598-0185
Fax: +1-317-598-0187

Or visit us on the web at www.doubletake.com



Get the standard today: www.doubletake.com or 888-674-9495

