

Realtime
publishers

Protecting Data with a Unified Platform

The Essentials Series

sponsored by



Dan Sullivan

Introduction to Realtime Publishers

by Don Jones, Series Editor

For several years now, Realtime has produced dozens and dozens of high-quality books that just happen to be delivered in electronic format—at no cost to you, the reader. We’ve made this unique publishing model work through the generous support and cooperation of our sponsors, who agree to bear each book’s production expenses for the benefit of our readers.

Although we’ve always offered our publications to you for free, don’t think for a moment that quality is anything less than our top priority. My job is to make sure that our books are as good as—and in most cases better than—any printed book that would cost you \$40 or more. Our electronic publishing model offers several advantages over printed books: You receive chapters literally as fast as our authors produce them (hence the “realtime” aspect of our model), and we can update chapters to reflect the latest changes in technology.

I want to point out that our books are by no means paid advertisements or white papers. We’re an independent publishing company, and an important aspect of my job is to make sure that our authors are free to voice their expertise and opinions without reservation or restriction. We maintain complete editorial control of our publications, and I’m proud that we’ve produced so many quality books over the past years.

I want to extend an invitation to visit us at <http://nexus.realtimepublishers.com>, especially if you’ve received this publication from a friend or colleague. We have a wide variety of additional books on a range of topics, and you’re sure to find something that’s of interest to you—and it won’t cost you a thing. We hope you’ll continue to come to Realtime for your educational needs far into the future.

Until then, enjoy.

Don Jones

Introduction to Realtime Publishers..... i

Data Protection in Dynamic Environments 1

 Business Drivers Behind Dynamic Environments 1

 Increasing Productivity..... 2

 Growing Importance of Intellectual Property and Business Data..... 2

 Controlling Costs..... 3

Distribution of Business Data Over Multiple Devices 3

Application-Specific Requirements for Backup and Recovery 4

Business and Technical Constraints on Data Protection Options 5

Summary 5

Copyright Statement

© 2012 Realtime Publishers. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers (the “Materials”) and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtime Publishers its web site sponsors. In no event shall Realtime Publishers or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtime Publishers and the Realtime Publishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at info@realtimepublishers.com.

Data Protection in Dynamic Environments

Businesses operate in a constantly changing environment. Demands for products and services shift over time. Organizations respond to the actions of competitors and collaborators. Innovation creates new opportunities and sometimes threatens existing ways of doing things. The impact of these changes is not limited to the confines of the executive suites filled with people trying to chart a course through the changing environment. The impact of these changes can permeate an entire organization. IT departments, in particular, are required to respond to these changes. Regardless of what else changes for IT professionals, there is always a need to protect data and ensure operations can continue.

Protecting against data loss in such a dynamic environment requires attention to multiple issues:

- Business drivers behind these dynamic environments
- Distribution of business data over multiple devices
- Application-specific requirements for backup and recovery
- Business and technical constraints on data protection options

This article will discuss these issues with particular attention to how they affect data protection strategies and methods.

Business Drivers Behind Dynamic Environments

The global economy is a complex system that is difficult, perhaps impossible, to understand in fine-grained detail. We can, however, identify broad patterns and characteristics that can affect virtually any business:

- Increase productivity
- Growing importance of intellectual property and business data
- Control costs

Although these high-level factors do not directly affect the way organizations protect data, they do have indirect consequences for IT operations.

Increasing Productivity

Productivity is a measure of value derived from investment in human and capital resources. One way a business can increase productivity is to improve the way it uses technology, especially information technology. Information technologies can reduce costs by:

- Replacing manual processes with automated procedures
- Improving information sharing among employees and reducing the costs of sharing that information
- Improving data analysis methods that provide information to business decision makers
- Introducing new products and services based on information technology that lead to new revenue streams for the business

There is, of course, a cost for all of these kinds of improvements, including the cost of backing up data and applications and ensuring they can be recovered if needed.

Growing Importance of Intellectual Property and Business Data

Businesses measure their success, in part, by their assets. Assets such as real estate, manufacturing equipment, and computers are clearly valuable. However, intangible assets, such as business processes, patents, trade secrets, and proprietary data, can be a substantial part of a company's assets.

One way to informally assess the value of your business data is to imagine you were to hand all that data over to your competitors. How valuable would it be for them to have your customer list, strategy documents, marketing plans, software designs, email messages, and databases? Now consider the possibility of losing some portion of that data. How would your business recover or reconstruct that data? What would be the cost of reconstructing the data if you did not have adequate backups? Questions such as these can help you understand the importance of intangible assets such as intellectual property and business data.

Business data is constantly changing in the course of normal day-to-day business. Innovations might not be as frequent but changes in business processes and slight improvements in software or product design accumulate to produce overall changes in a business' intellectual property.

Controlling Costs

Businesses constantly seek to control costs, and information technology is one means of doing that. New software may be deployed to automate a formerly manual process. A new collaboration tool may be rolled out to reduce the time needed to resolve customer support issues. A business intelligence portal might be deployed to give line managers better information about their operations so that they can more effectively control costs. These kinds of changes will require changes in the IT environment, and that means there will be changes to the way data is backed up and restored as well. At the same time these business drivers are reshaping the IT environment, technical changes are altering where business data is stored.

Distribution of Business Data Over Multiple Devices

Business data is rarely found in a single monolithic database or even on a single server. Data tends to be structured to align with business functions and application requirements. Financial data may be stored in a centralized back-office application database while sales information, proposals, and quotes may be spread across a large number of desktops and laptops.




Figure 1: Business data is distributed over a wide range of devices including employee-owned mobile devices and public cloud providers.

In addition to traditional business-owned servers, desktops, and laptops, there is increasing use of employee-owned devices for business purposes. Organizations should consider the pros and cons of bring your own device (BYOD) practices and device policies that allow for the benefits of BYOD for both employees and employers while mitigating the disadvantages. Part of those policies should include consideration of what, if any, business data is allowed to be stored on employee devices, the use of encryption if needed, and the need for backups of data stored on mobile devices.

The distribution of data over multiple devices is even more pronounced when you consider virtualization. With virtual machines, business data might be stored under different OSs even though it may reside on the same storage device. At the lower levels of implementation, all data on a device is stored in a consistent format but backup operations typically occur at the operating system and hypervisor levels. Different hypervisors will implement different services for backup programs, such as identifying changed blocks to minimize the number of blocks that need to be copied during a backup operation. This reality is just one of the ways hypervisors introduce additional complexity to the backup process.

Virtualization is so easily implemented with today's hypervisors that virtual machine sprawl can become a problem. If you are not careful to track when and how virtual machines are created, you might create virtual machines that store business data that is not included in backup schedules, leaving business data unprotected. In addition, some of the most frequently changing data repositories in an organization have their own specific requirements for backup and recovery.

Application-Specific Requirements for Backup and Recovery

Office documents, such as word processing documents and spreadsheets, are easy to back up. Files used for databases, email servers, and collaboration servers are much more difficult to back up. The problem stems from the way files are used in these systems.

Consider how a basic relational database works. There are multiple schemas, which are collections of tables and other database objects related to a single subject area, such as product inventory or customer contact information. Tables in these schemas consist of rows of data such as customer names, addresses, and contact information. There are also ancillary data structures such as indexes that are stored within the database. Rather than implement separate files for all these schemas, tables, and indexes, backups often use a set of files that each store multiple data structures. The files can become large; large enough that it is not efficient to make an entire backup copy when there are only some changes to the data.

This kind of application-specific data storage mechanism allows databases, email servers, and collaboration servers to function efficiently. They impose constraints on backup systems, though. Efficient backup programs should accommodate the application-specific needs of these programs.

Business and Technical Constraints on Data Protection Options

Given the requirements for data protection that have been outlined here, it is important to consider the constraints on data protection imposed by business considerations as well as technical limitations.

Storage costs can be a significant factor in backup and recovery operations. Inefficient backup operations and lack of data compression can lead to less efficient storage use.

Consider how data must be moved in order to perform backups. Large volumes of data can generate substantial traffic on a network. Care should be taken when backing up virtual machines so that multiple virtual machines are not contending for the same physical devices, such as network interfaces, at the same time.

Scheduling is another factor to consider when creating enterprise backup procedures. Backups should be scheduled frequently enough to provide adequate recovery points. For example, if backups are only performed once every two days, you could potentially lose up to two-days-worth of work. Critical applications can warrant more frequent backups to reduce the risk of losing data.

Policy-driven backup procedures can help ensure backups are performed in a consistent manner across multiple devices. Policies that replace device or application-specific backup scripts can help reduce the possibility of an undetected error in any one of a large number of scripts.

Summary

Business is dynamic and thus IT operations are subject to frequent change. A number of different drivers influence how businesses operate; some are organizational issues and some are technical. Changes in the way businesses use technologies, such as mobile devices and cloud computing, have a direct impact on data protection. Business factors, such as changing market conditions and industry innovation, can have more indirect impact on data protection. Regardless of the source or motivation for changes, IT professionals still face the constant requirement to preserve essential business data.