Realtime publishers

The Essentials Series: The Business Benefits of Rapid Application Failover

Defining the Peril of Application Downtime

sponsored by

ARCserve® More than Backup

by Greg Shields

Defining the Peril of Application Downtime 1
Applications Are as Critical as Data 2
Surviving the Interruption 3
Replication Technology as Downtime Eliminator 4

Copyright Statement

© 2009 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 or 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.

Defining the Peril of Application Downtime

You're a smart IT professional. You manage your servers, network, and applications with the right level of due diligence. You respond to customer requests within acceptable timelines. You always back up your data. You've probably even negotiated a Service Level Agreement (SLA) with your business, codifying the maximum acceptable level of downtime and setting the standard for which services need added attention.

You've done everything right, and as a result, you've protected your business from the threat of data loss. That's an important milestone, but it really solves only one part of the problem. Ensuring the safety of your business' data protects its intellectual property, *but that data is effectively worthless without the applications that make use of it.*

A business' worst-case scenario often arrives through a data center-impacting incident. The cause for that incident can range from a simple disruption to a major outage on up through the catastrophic losses caused through a natural disaster. Yet while the damage done by an interruption in data center operations ranges in severity, the net result to your users is the same: They can't do their jobs.

Back in 2006, the National Academy of Archives and Records reported that 96% of companies who lose access to their data centers for a period of 10 days or longer are out of business within a year (Source: *Continuity Insights:* http://www.continuityinsights.com/Magazine/Issue_Archives/2006/05-06/Virtually_Painless.html). This statistic should strike fear into the heart of any data center manager or IT director. Proactive risk management should recognize that even the simplest of non-destructive outages—extended power loss, minor water damage, air conditioning unit failures, and so on —can very quickly turn into a business-critical incident.

A further study by McGladrey and Pullen showed that 43% of companies who experience a disaster will never recover. The major functional difference between those that survive and those that don't has to do with how fast their IT can swing itself back into operations. In short, your planning must prepare and compensate for all types of interruptions, from the outwardly benign to the completely devastating.

Applications Are as Critical as Data

Consider for a moment the types of data in your organization today. Some data is generally processed through client applications; for example, a Microsoft Word document requires Microsoft Word to view or edit its contents. Other types of data require no specialized application whatsoever for their processing, being able to work with essentially any application on any desktop or server. Other types of data require the presence of server resources for their processing, such as a server application, a database, or both.

From a high-level perspective, these different types of data can be functionally broken down into four categories:

- **Raw data.** This data requires no specified application to view or edit its contents. This kind of data is typically composed of raw text files. Although this unstructured data has the greatest level of malleability in your computing environment, most IT environments tend to store relatively little of it.
- **Desktop-centric data.** Other types of data require applications at the desktop for their processing. These desktop-centric applications, such as Microsoft Office or a single-user homegrown application, load data into the desktop application for viewing and editing. That data is later saved onto traditional file servers where it can be accessed by others. Due to the desktop-centric nature of this kind of data, neither this data nor its applications are shared.
- **Server-centric data.** Other kinds of data require the support of server-based applications for their processing. For this kind of data, a network connection to the server and its hosted application is required. Some examples of applications that generate and use this kind of data are Siebel CRM solutions or Microsoft Exchange email. Because there is a direct linkage between the data and the application that processes that data, the functionality of the application is required for the data to be processed.
- **Database-centric data.** The last form of data relates to that which is stored in your business' databases. Information here can be viewed and manipulated through desktop-centric applications, or it can require a server-based application, or both. For this reason, database data often requires the functionality of both the database and the processing application. For example, the same data that leverages a Siebel CRM application can actually be stored in an Oracle database. Alternatively, remote access to Microsoft Outlook can require services from both the remote access server and the Microsoft Exchange server.

The effective IT organization recognizes that this separation in data exists. Such organizations are responsible for its safety and security as well as its non-interrupted availability. With backup and restore solutions being relatively mature—tape backups, disk-to-disk solutions, continuous data protection technologies—it can be argued that essentially all IT organizations do a good job with data backups.

Yet herein lies the problem: *Most IT organizations focus on data as if all data were raw.* IT organizations with solid backups and offsite replication can and do successfully ensure their data will survive through disasters both large and small; however, no backup solution alone can fully preserve that data's *accessibility* without also considering its needed applications.

Surviving the Interruption

Does your organization currently consider the role of applications in its planning for disaster recovery? How about for comparatively minor service interruptions? If you have an SLA in place that codifies your IT organization's Recovery Time Objectives (RTOs), is that number scoped to include your users' ability to process the data?

The peril of application downtime relates strongly to the capability for your IT organization to maintain a functional percentage of total services during the period of a service interruption. Most organizations recognize that keeping every service up and operational isn't feasible when their "major outage" light turns from green to red. Most recognize that key businesses services must be brought back online quickly for the business to continue as a business.

Yet there's a technological problem native to applications and their preservation during service interruptions. Unlike data, which tends to be static while at rest, the applications that process data are constantly processing and constantly changing. Thus, traditional backup solutions must evolve to support the real-time needs of applications and their internal data processing. These solutions must work with each server's individual operating system (OS) to recognize when changes have been made and replicate those changes to a backup host.

Further, today's business needs mean that today's technologies must evolve past the idea of "tape drive as primary backup solution." The tape backups of yesteryear were successful in gathering an emergency copy of critical data to protect against accidental loss and ensure long-term archival. But large-scale tape restorations such as the kind that occur after a major service interruption or disaster can require hours, if not days, to complete. If the loss of a single terabyte of critical data requires 3 days to restore, this solution involves a risk to your business operations.

Replication Technology as Downtime Eliminator

As a result, alternative technologies are necessary that enable the rapid and real-time failover of business-critical data to alternative locations. At the same time, these same solutions enable the failover of that data's applications. Technologies exist today that enable the smart business to inexpensively and cohesively failover every aspect of their data processing to alternative sites. The result is the elimination of downtime associated with all kinds of interruptive events, from simple power outages all the way to catastrophic disasters. In the next article of this series, I'll focus on just those solutions, helping you understand exactly how these solutions work and where they might fit into your overall data protection portfolio.