

Realtime  
publishers

*The Definitive Guide<sup>™</sup> To*

# Windows Application and Server Backup 2.0

*sponsored by*



*Don Jones*

Chapter 11: Redesigning Backup: Is It Really Worth It?.....	181
Let's Review: What is Backup 2.0? .....	181
A Backup 2.0 Shopping List.....	187
What's Involved in a Redesign? .....	188
Calculating the Real Cost of Your Grandfather's Backup Strategy .....	191
Technical and Business Advantages of Backup 2.0 .....	193
Whole Servers .....	193
Exchange.....	194
SQL Server.....	194
SharePoint.....	194
Virtualization .....	194
Is a Redesign That Expensive? .....	194
Tackling the Office Politics .....	195
Conclusion: This Is Not Your Father's Backup Plan.....	196
Coming Up Next... ..	197

## **Copyright Statement**

© 2010 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](mailto:info@realtimepublishers.com).

**[Editor's Note:** This book was downloaded from Realtime Nexus—The Digital Library for IT Professionals. All leading technology books from Realtime Publishers can be found at <http://nexus.realtimepublishers.com>.]

## Chapter 11: Redesigning Backup: Is It Really Worth It?

---

What's it going to cost you to implement Backup 2.0, and is it worth the time and money? Obviously, I can't give you specific numbers because those numbers will depend on exactly what you're backing up, how distributed your network is, and a number of other factors. But what I *can* do is show you how to calculate that cost, and to calculate the cost of just staying with your existing backup infrastructure. You can do the math and figure out whether a 2.0-inspired redesign is going to be beneficial for you.


### Let's Review: What is Backup 2.0?

Before we do that, let's quickly review what Backup 2.0 is all about. As always, we're focused on a business goal here, and we're not concerning ourselves with past techniques or technologies. Here's what we want:

Backups should prevent us from losing any data or losing any work, and ensure that we always have access to our data with as little downtime as possible.

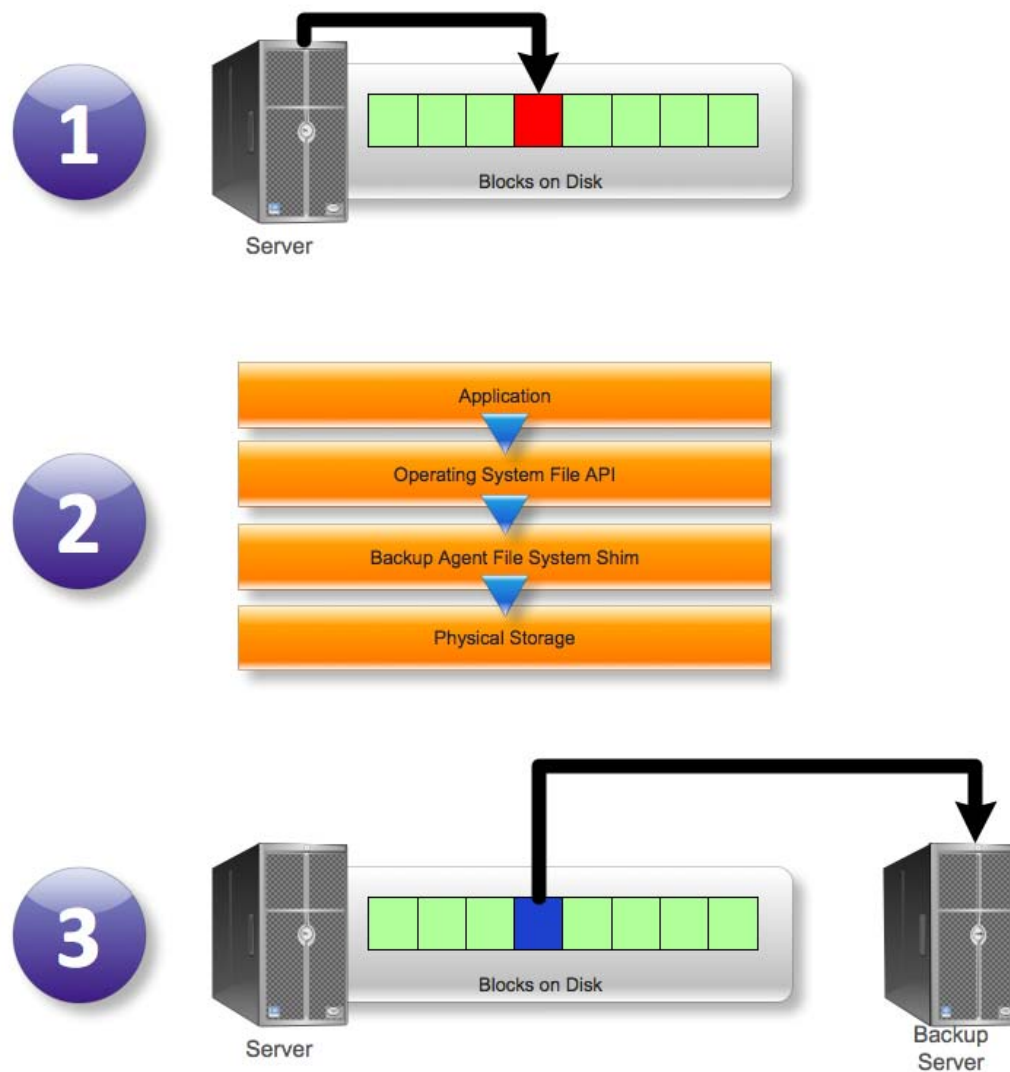
The problem with old-school backup techniques is that they rely primarily on point-in-time snapshots and on relatively slow tape-based media for primary backup storage. That means recovery is always slower than it should be, and we always have a lot of data at-risk.

Backup 2.0 proposes that we use *continuous data protection*. In a typical implementation, you would install agent software on all your servers, and they would back up your data to a central backup server. Figure 11.1 illustrates this arrangement. The backup server would keep your backup information on disk, making it more readily available. With a fast 10Gbps Ethernet network connection, you could restore an entire server in a few minutes—far faster than you could from even the fastest tape drive.



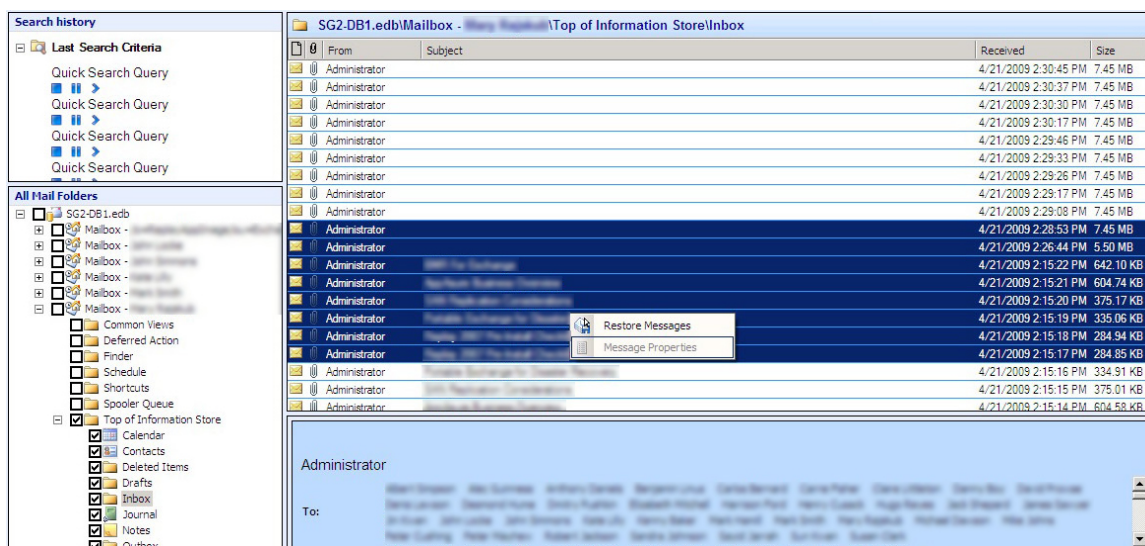
**Figure 11.1: Backing up to a central, disk-based backup server.**

Rather than grabbing snapshots of the entire server or even entire files, the agent would “shim” the operating system’s (OS’s) file system, enabling it to “see” every change being written to disk at a disk block level. Because the agent would be accessing this data below the level of a file, it wouldn’t have to worry about “open” files or any other common backup restrictions; it could simply copy every disk block change as it was being written to disk. Figure 11.2 shows how this might work. In Step 1, an application makes a change to disk—perhaps saving a row to a database, writing a new message to a message store, or modifying a file. In Step 2, those changed disk blocks are captured by the file system shim, and in Step 3, they are transmitted to a backup server, which saves the disk blocks in a database.




**Figure 11.2: Backing up disk blocks.**

Because the disk blocks are time stamped, the backup server could re-construct the original server's entire disk at any given point in time. By exposing that information as a mountable volume, and by providing utilities that understand how to navigate databases such as SQL Server or Exchange databases, you could retrieve single items from any point in time. Figure 11.3 shows how a user interface might expose that kind of functionality.

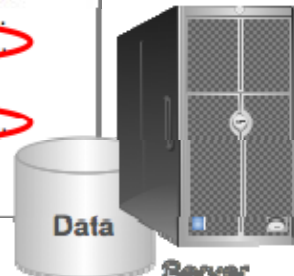


**Figure 11.3: Recovering single items from the backup database.**

To reduce the size of the backup database, the backup server could use file- or block-level de-duplication along with compression. For example, as illustrated in Figure 11.4, the server could detect disk blocks that contained the same information and, as shown in Figure 11.5, remove the duplicates and simply provide a note of what block they were a duplicate of. Vendors implementing this technique, along with compression, claim up to 80% reductions in data size.

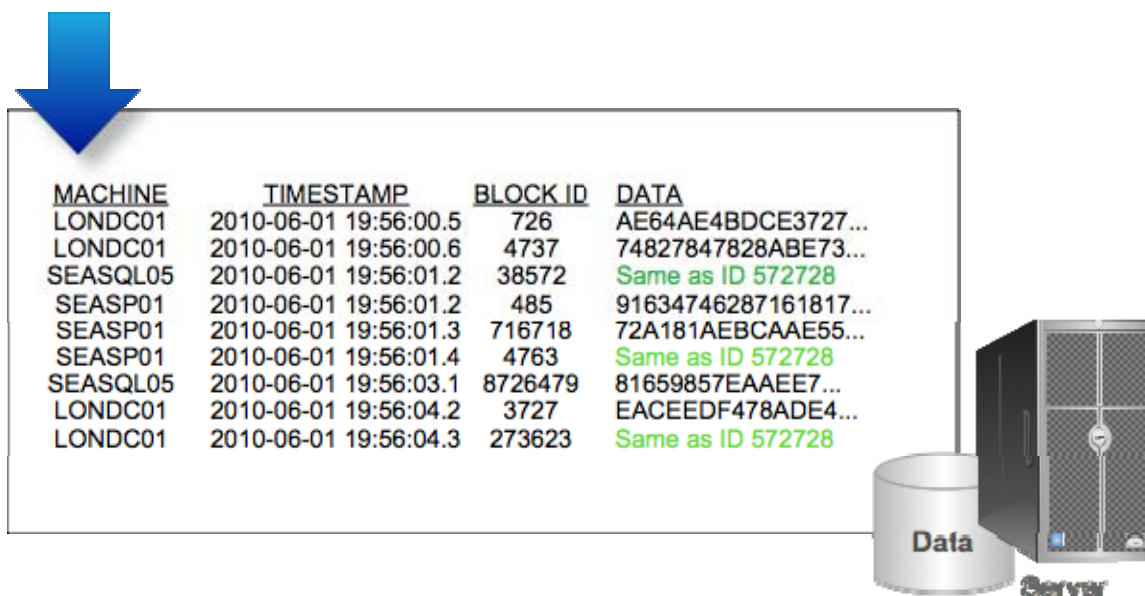


MACHINE	TIMESTAMP	BLOCK ID	DATA
LONDC01	2010-06-01 19:56:00.5	726	AE64AE4BDCE3727...
LONDC01	2010-06-01 19:56:00.6	4737	74827847828ABE73...
SEASQL05	2010-06-01 19:56:01.2	38572	AE64AE4BDCE3727...
SEASP01	2010-06-01 19:56:01.2	485	91634746287161817...
SEASP01	2010-06-01 19:56:01.3	716718	72A181AEBCAAE55...
SEASP01	2010-06-01 19:56:01.4	4763	AE64AE4BDCE3727...
SEASQL05	2010-06-01 19:56:03.1	8726479	81659857EAAEE7...
LONDC01	2010-06-01 19:56:04.2	3727	EACFEDE478ADE4...
LONDC01	2010-06-01 19:56:04.3	273623	AE64AE4BDCE3727...




**Figure 11.4: Detecting duplicate disk blocks...**





**Figure 11.5: ...and de-duplicating them to save space.**


Magnetic tape wouldn't necessarily be eliminated. Tape still provides a convenient way to move data safely off-site. However, tapes would be used to back up the *backup database*, rather than servers directly, as Figure 11.6 shows. This would effectively eliminate backup windows: You could run tape backups all day long, if you wanted to, because you wouldn't be impacting your servers.



**Figure 11.6: Moving backups off-site via tape.**



However, because Backup 2.0 deals in disk blocks that are captured in near-real time, you could also stream those disk blocks off-site to a second backup server by means of a WAN or Internet connection. “Off-site” might refer to another facility of your own or to a contracted disaster recovery facility. Figure 11.7 illustrates this technique; it can be very useful in creating a “warm recovery site” that can begin restoring key servers to virtual machines in the event of a total disaster in your main data center.



**Figure 11.7: Replicating backup content off-site by means of a WAN or Internet connection.**

These are the basic capabilities that Backup 2.0 offers, and the core techniques that it uses. It’s designed to provide both rapid single-item recovery as well as speedy whole-server recovery, giving you the flexibility of restoring entire servers to the same hardware, to different hardware, or even to a virtual machine.

### Let's Be Clear

The capabilities I've written about here are largely wishful thinking on my part, meaning that I'm not writing about any specific vendor or solution. That's deliberate: I don't think we should acquire technology based on what's available. Rather, I think we should define what we want, then push the marketplace to deliver.

There are absolutely vendors who offer at least some, if not all, of the capabilities I've written about. Some of them use the phrase "Backup 2.0" and some don't. The key for you as a technology purchaser is to focus on the *capabilities your business needs*. When you find a vendor who can deliver exactly those capabilities for the price you feel they're worth, you've found your new backup solution.

## A Backup 2.0 Shopping List

Let me quickly summarize the key capabilities of a Backup 2.0 solution:

- **Continuous data protection** is the key element, leaving little or no data at-risk in the event of a failure of any kind. There are many ways in which to implement this, but my preference is to focus on streaming changed disk blocks. Doing so, rather than grabbing entire files, bypasses the need to worry about files that are constantly open and provides more granular recovery capabilities.
- **Mountable recovery images** are another key element. You should essentially be able to browse your backups as if they were a file system. Depending on your environment, this might include the ability to mount specific kinds of databases—such as SQL Server, SharePoint, or Exchange—so that you can recover single items without having to restore the entire database.
- **Variable recovery targets** offer the ability to recover a server to dissimilar hardware or even to a virtual machine. This capability provides maximum flexibility for recovery scenarios.
- **Data compression and de-duplication** reduce the amount of disk space consumed by the backup solution.
- **Tiered storage** allows you to transfer or copy backup data to tape for off-site storage.
- **Replication** allows you to transfer or copy backup data to another backup server for redundancy, off-site storage, centralization, or other goals.


Your backup solution has to provide these capabilities, plus any others that you've identified as important to your business. Redesigning your backup infrastructure means acquiring a compliant solution, then redesigning—if necessary—your data center to support that solution.

## What's Involved in a Redesign?

If you're like most companies, you probably already have a backup plan in place. That most likely involves one of three basic scenarios:

- Backup tape drives hanging directly off servers, which back up to that drive
- Backup tape libraries hanging off a dedicated backup server, which uses software agents to grab snapshots and stream them to the server, which then streams the data to tape
- A combination of these two, as Figure 11.8 shows

The third option is actually fairly common; companies start with the centralized tape backup scheme, then attach tape drives directly to a few key servers to enable more backups to take place within a given backup window.



**Figure 11.8: Typical tape-based backup.**

My first IT job was as an AS/400 operator, and my main job was backing up the system every night (to this day, I hate night shifts). I remember my elation when we purchased a new tape backup drive that was capable of backing up the whole system in half a shift, using just one rack of tapes. It replaced a system that took almost the whole night and two racks of tapes. That's when I first started to think, "Wow, our expectations for backup and recovery are awfully low if this is what excites me." That was the beginning of the "Backup 2.0" concept for me: Backups that actually meet business requirements rather than fit within technology limitations.

It's not even unusual to have a sort of hybrid model, where the backup server collects backup information on disk, then writes it to tape, thus providing more backup capacity in a shorter timeframe. The backup server typically doesn't store more than a day or two of snapshots, but that's often enough to provide recovery to the most-recent image.

Redesigning this scheme to fit the Backup 2.0 model isn't difficult. You're obviously going to need to acquire a Backup 2.0 continuous data protection solution; feed "continuous data protection" to your favorite search engine and see what comes up. You can also just try searching on "Backup 2.0," as several vendors are starting to use that term to describe their products.


#### Caution

Please don't assume that a vendor's use of the term "Backup 2.0" guarantees that they deliver every single capability I've written about. I can't control the marketing departments of every vendor out there; use this book to develop your own list of criteria and requirements, and review each solution that you're considering accordingly.

You *will* have to "touch" all your servers, or use some central software deployment mechanism: You're going to want to uninstall or disable your old Backup 1.0-style agents on each server, and install agents for your new Backup 2.0 solution. In all likelihood, you can continue using your existing backup server, with its attached tape library. You'll want to change any servers that have directly-attached tape drives and have them instead stream disk blocks to your backup server. Figure 11.9 shows the new arrangement; at worst, you might need to add disk storage to your backup server.

#### Note

Frankly, determining the amount of disk storage your "central backup server" will need is probably the biggest challenge in moving to Backup 2.0. *Every single vendor is different* in this space; work with them to determine some disk space estimates. Just remember that your plan will probably be to move backup data to tape on some scheduled basis, and remember that your backup system isn't necessarily supposed to be a *data archival* system. You don't need to keep everything forever; you need to keep enough to bring your business back online in the event of a disaster or to recover corrupted or accidentally-deleted files.



**Figure 11.9: Converting to Backup 2.0.**

Ask your Backup 2.0 solution vendor if they offer compression and de-duplication, and for some realistic expectations on how much space that will save. Let's say it's 60%. Do some analysis to see how much disk space actually changes on your servers on a daily basis. Decide how many days into the past you want to keep on live, disk-based storage. Then do the math:

- Multiply your daily disk change by the number of days you want to keep
- Multiply that result by the average disk reduction number (say, .7 for 70%)

Add a 10% pad for growth and mis-measurement, and you should have an idea of how much disk space your backup server will need.

You also need to look at the network your backup server will be running over. You might have already built a dedicated network in your data center for backup traffic; if so, it will probably be just fine in the new Backup 2.0 world. If not, you might want to consider doing so—and consider a fast 10Gbps Ethernet network for that purpose. Remember, this is the big change from Backup 1.0: You're not using that backup network solely at night during your maintenance window; you're using it *continually*, so it needs to be fast, and ideally it needs to be dedicated.

### Note

Your investment in a fast backup network is *not* to enable fast backups, contrary as that might sound. The investment is to enable faster *recovery*. Remember, with continuous data protection, your backups are always up-to-date; the goal is to be able to get them *off* the backup system and onto your actual servers as quickly as possible. That's why a 10Gbps Ethernet network dedicated for backups is a great idea. It's not always a *necessity*, of course, but it's great if you can afford it.

Add up your costs:

- Your Backup 2.0 solution, which is likely going to be licensed based on the number of servers it will back up
- New disk storage for your backup server, if needed—This can be pricey; ideally, you want a storage area network (SAN) for this purpose, although iSCSI makes a SAN less expensive than older fibre channel SAN technology
- If you need to implement a dedicated backup network, factor in the price of network cards (if your servers don't have unused Ethernet ports), cabling, switches, and other infrastructure equipment

Those are numbers you should be able to come up with pretty readily, given a quote from a few vendors and perhaps a trip to a search engine to check prices on the networking equipment or disk storage. Next, you need to figure out what it costs you to *do nothing*. Why? I don't want you changing your whole backup infrastructure on my say-so (I know you wouldn't anyway). It should be a business decision, based on what you'll gain from the investment.

## Calculating the Real Cost of Your Grandfather's Backup Strategy

Determining the cost of "your grandfather's" backup strategy, unfortunately, is a tough thing to work out because you're dealing with risk numbers, not hard-and-fast pricing. Ask any company "How much money do we lose per hour if such-and-such a server is offline" and you probably won't get a very accurate answer; most businesses simply don't know.

But let's look at the risks associated with doing nothing.

In a simple restore scenario, costs are actually a bit easier to figure out. Let's say someone asks you to recover a file from a week ago, that you have the necessary tapes in hand, and that the recovery information is only on those tapes—it's not sitting on disk someplace. How long, on average, will it take someone to completely restore that file? Figure that out, and then figure out how much both that person *and* the person waiting on the file are being paid for that period of time. Now, multiply by the number of times you have to do this in a year, and that's what Backup 1.0 costs.

### Note

If you don't have a way of tracking how often this is done in a year, you should get some kind of Help desk ticketing software so that you *can* track it. That should also let you come up with an average time-to-restore. If you aren't tracking that information, it becomes very difficult to make a case for improving the situation.

I'll give you an example from my own past: It took us, on average, 30 minutes to recover a single file. We had to pull tapes out of our safe, load them, locate the file in the backup software's index, then let it spin the tape to the right location to retrieve the data. Frankly, 30 minutes on average is pretty good. Our administrator made about \$60,000 a year at the time, and the folks in the office probably made a similar amount, on average. I recall one year where we did at least three of these operations a week—so about 150 a year. There are about 1,920 working hours in a year, and these backup operations consumed about 75 of them. That means this kind of recovery operation cost us about \$4700 per year. I know that's not a huge number—a single server usually costs more—but it adds up. And we were a relatively small company; \$5k a year in overhead was a healthy percentage of our IT budget.

### Note

Features like Windows' Previous Versions (Volume Shadow Copy) were implemented specifically to help reduce that cost. By giving users a self-service mechanism and by using faster disk-based storage for recent backups, you could probably reduce that cost by a lot. But Previous Versions has some significant limitations, including its snapshot-based methodology and the limited disk space most companies can give it.

Calculating the cost of an entire server disaster—or worse, a data center disaster—is obviously much more complicated. How many users are impacted? How much of their productivity is lost? Are customers directly impacted? If so, how much in sales will you miss while the server is down?

Consider a database server that serves as the back-end for an e-commerce application. If the database is lost through corruption or failure, or you need to restore the entire server for some reason, how much in sales does the company lose per hour? In most cases, a single hour of lost sales can justify even the most expensive backup infrastructure redesign.

You also have to consider some of the major technical and business advantages of Backup 2.0. These can't always be measured directly in monetary terms, but they do offer flexibility and improvements that a business can appreciate.



Want to know what I like to do? Since this whole “how much does a failure really cost us” question is so hard to really nail down, I like to have a little conference with the company executives. “How much are you willing to spend,” I’ll ask, “to make it so that I can have your data center back online within a few hours of it burning down?” That kind of tells you how much the shareholders are willing to spend on continuous data protection. The second question is, “Now, when I bring you back up, how much work—in the middle of a disaster, mind you—are you willing to re-do after I get the system back online? How much money are you willing to spend to make that answer ‘none?’”

I’ve often been shocked by the answers. I once estimated that a consulting client of mine would lose \$10,000 per hour if they were totally offline. I put together a Backup 2.0-style recovery solution that would have cost about \$30,000. Easy math, I thought. I asked those two questions of the CEO, and he said, “I’d spend \$250k a year to get back online within an hour, and another \$250k a year to ensure I didn’t lose more than 10 minutes’ of work when we came back online.” Well, then. I was clearly under-pricing myself! Mind you, he didn’t have any hard numbers—that was just a gut instinct from the person who had to answer to the owners—and sometimes, that’s a better number than you might think.

## Technical and Business Advantages of Backup 2.0

As I’ve written, sometimes it can be difficult to really quantify, in dollars, the cost associated with *not* moving to a Backup 2.0-style approach. What can sometimes help is to understand some of the specific business and technology advantages of this new approach so that you can perhaps quantify those *advantages*. Doing so can help you justify the cost of a redesign in terms of “benefit gained” rather than “risk avoided.”

### Whole Servers

Backup 2.0, at its heart, offers the ability to finally grab a solid backup of every single server in your environment, enabling you to recover anything from a whole server to a single file, to a specific point of time, within seconds. Literally *seconds*, especially with individual files or groups of files; perhaps minutes, depending on the infrastructure you build, for whole servers.

And when it comes to whole-server recovery, Backup 2.0 supports a vast array of new options: Recover to a virtual data center, either on-site or off-site. Recover to dissimilar hardware. Recover to the same hardware. Whatever you want—you can flex these options as needed for any particular scenario.

## Exchange

Backup 2.0 offers a far superior approach for Exchange Server, which is perhaps one of the most mission-critical parts of any business' infrastructure these days. Recover individual messages *without having to restore an entire database*, and without having to stream those individual messages out your database one at a time the way Backup 1.0 does. Backup 2.0 gives you the ability to search, find, and recover messages without impacting the production Exchange Server, while still giving you the ability to recover entire servers with a few button clicks—either to physical or virtual hardware, giving you the flexibility to quickly come back online after almost any kind of failure.

## SQL Server

SQL Server has always been a bit easier to deal with in terms of recovering databases, but Backup 1.0 still (usually) requires you to restore an entire database to a SQL Server someplace—time-consuming and potentially impactful on production servers. Backup 2.0 should let you access backed-up databases—down to individual tables, even—without actually restoring the database to a production server. As with Exchange, this granular, low-impact ability to recover anything from a single table to an entire server offers you flexibility, speed, and efficiency.

## SharePoint

As SharePoint grows more and more embedded in our organizations—taking a role not unlike Exchange's in terms of its criticality to our businesses—the ability to granularly recover specific items, or restore an entire SharePoint server or farm, becomes ever more important. Backup 2.0 can meet the challenge, provided you select a solution that's specifically designed to support this kind of SharePoint backup and recovery.

## Virtualization

Finally, Backup 2.0 is perhaps the first time we've been able to properly address our growing virtual infrastructure. Treat virtual machines as you would a real one, with low-overhead, block-level backups to a centralized server. You're essentially able to take virtualization "out of the mix," treating virtual servers exactly the same as you would physical ones.

## Is a Redesign That Expensive?

I've found that redesigning for Backup 2.0 is typically *not* that expensive. Let's take a worst-case scenario, which in my consulting work typically means having to

- Acquire a Backup 2.0 software solution—often the most expensive part of the proposal
- Implement a back-end "backup network" in the data center to carry the backup solution's traffic
- Perhaps upgrade or acquire tape libraries to provide a second tier of backup storage

In a typical small- to midsize business, these costs are often small enough to be considered incremental. It's impossible for me to provide accurate price quotes because every business will be slightly different, but rest assured that the cost is typically in-line with implementing a more traditional backup solution. In other words, figure out what your existing Backup 1.0 solution costs you, and you're probably looking at similar costs to implement Backup 2.0. Often less, as you'll usually be able to re-use your existing backup server(s) and tape drive(s), and if you already have a back-end network for backup traffic, you'll be able to re-use that as well. In a best-case scenario, in fact, you're really just swapping your current backup software for new backup software, and that's rarely a massive investment.

Perhaps the biggest concern I run across is with the company's CFO because computer software is often depreciated over a longer period of time than even computer hardware. In other words, if your CFO doesn't feel he's depreciated your old backup software, there may be some resistance to buying new software. That's where the "pros and cons" side of the argument comes in: You might need to identify sufficient benefit that losing the depreciation on the old software is justified. Also be sure to factor in any software maintenance fees that you pay on the old software and that you'll pay on the new software; because those maintenance fees are recurring, they're typically not depreciated at all—and those fees can sometimes constitute the bulk of the cost of a software product over a given period of time.

## **Tackling the Office Politics**

The final thing you'll probably have to deal with are the inevitable office politics that surround any kind of major IT change. The key to success in this arena is to be prepared with facts and to ask that everyone involved try to avoid becoming "religious" about the issue.

It can be difficult: We technology folks get attached to our toys, and we don't really like change very much (although we enjoy inflicting it on others when it's something we're excited about). The goal should be to have everyone agree to focus on the business needs, and on how those needs can be best achieved.

I find that there are usually just a few categories of opposition to changing a company's backup strategy. Executive-level opposition is usually focused on cost, often with the added twist related to the depreciation of your existing solution—which I've already written about. Executives just need to understand the cost/benefit argument to decide whether the purchase is worth it. This book should help you be able to articulate the benefits, but don't *argue* for a new solution. Instead, present the benefits and the costs in an unbiased, businesslike fashion—you'll earn more respect with the executives and help them really decide if Backup 2.0 is worth it.

Middle-managerial opposition usually comes from a desire to avoid change. That's understandable, as change is what causes problems in most IT environments. The key here is to have a solid plan for implementation, and to make it clear that you don't want to proceed without such a plan in place. That may involve running old and new solutions in parallel in a limited pilot for some period of time so that you can understand and mitigate or address any problems that arise. Work within whatever change management framework your company uses—following that process will help relieve managerial concerns in most cases.

There's often a feeling that swapping backup schemes will be disruptive, but it doesn't need to be. Backup 2.0 can live side-by-side with Backup 1.0 in a phased, gradual approach that avoids disruption and offers time for familiarization.

IT professional-level opposition—that is, opposition from the folks who will implement and operate the Backup 2.0 solution—is usually simple fear or dislike of change. Everyone will have to learn a new set of tools, new techniques, and new ways of thinking. Mitigate that fear by clearly communicating the benefits of Backup 2.0. Typically, I find that most administrators aren't thrilled with the overhead and tedium of their existing backup solution, and the benefits of Backup 2.0 from that level—once clearly explained—helps overcome resistance pretty quickly.

## Conclusion: This Is Not Your Father's Backup Plan

You know, the real thing to remember here is that Backup 2.0 is new. That's why I keep calling it "2.0," much as I know that will lead to the term being overused by marketing folks (remember "Web 2.0?" This isn't the same). Backup 2.0 is focused on a whole new approach to backup and recovery—an approach that, frankly, wasn't even practical or feasible two decades ago when our current backup practices were created.

The problem with old-school backups is simply that they were design around the constraints of technology. Many of those constraints have been solved, at this point, while Backup 1.0 has more or less failed to keep up. Backup 2.0 takes an entirely fresh approach, with an acknowledgement of the problems of Backup 1.0. Rather than bolting on solutions—like "open file managers" and other workarounds—to fix Backup 1.0's problems, Backup 2.0 starts from the ground up with a new way of achieving *our actual business goals*.

In the end, that's what Backup 2.0 is really all about: *meeting business goals*. I think you should very carefully consider your existing backup solution's ability to meet your real business goals, and perhaps re-state those goals with a "this is what we'd have in a perfect world" approach. Then start evaluating newer, Backup 2.0-style solutions and determine whether they can't do a better job of helping you meet those newer, more real-world goals that you've stated. Let the business drive the technology, not the other way around.

## Coming Up Next...

Well, there you have it: Backup 2.0. I don't want to just leave you with a list of capabilities and a solid argument for the boss, though; I also want to give you some confidence that this stuff really works in the real world. So in the final chapter, which is coming up next, I'm going to share some "tales from the trenches"—real life stories of administrators who are living with Backup 2.0 right now. It won't all be sunshine and rainbows, though: You'll read some gritty stories about implementations that weren't well thought out, and other less-than-perfect situations. Why? Well, those are important lessons to learn. Hopefully, you can take those stories and ensure that your life with Backup 2.0 is smooth sailing.

## Download Additional Books from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this book to be informative, we encourage you to download more of our industry-leading technology books and video guides at Realtime Nexus. Please visit <http://nexus.realtimepublishers.com>.