



Strategies for Accelerated Backup and Immediate Recovery: Going Beyond Traditional Backup Methods

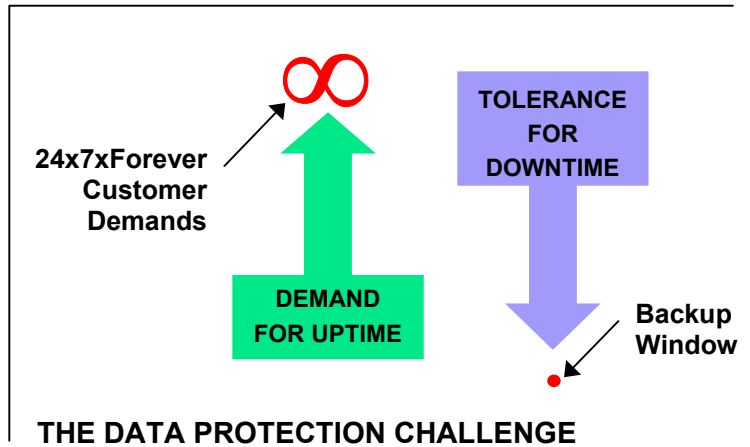
A WHITE PAPER

Abstract: In a world where the backup window has shrunk to virtually nil, and downtime is not an option, three things are crucial to enterprise survival: super-fast non-impact online backup; rapid, if not immediate, recovery, and a comprehensive backup plan that addresses not just one type of data loss, but all types. This paper explains how traditional backup methods have fallen short and how new methods achieve these ends for true data protection and surefire business continuity.

New Methods Solve Persistent Backup Challenges

The Data Protection Challenge

As data piles up exponentially on a daily basis, it's a multifaceted challenge that arises: it gets ever-more difficult to back up the mushrooming Terabytes of data within a rapidly shrinking backup window; it becomes harder to stomach long recovery times; transactional integrity grows imperative; it becomes increasingly likely that something *will* go wrong; scaling backup solutions gets very expensive while budgets get tighter; and, at the same time, the consequences of losing all—or any part—of that data keep increasing in severity. It's the bad dream that just keeps getting worse!

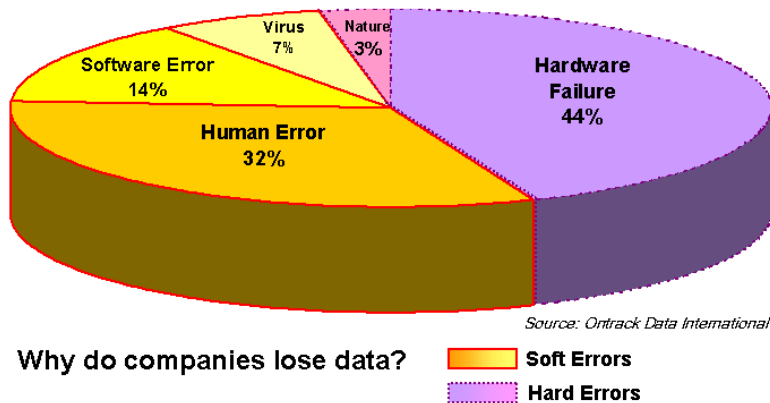


Limitations of Past Backup Paradigms

Up until recently, the inefficient and expensive traditional paradigm of each server backing up its own data to a local tape drive had progressively evolved into technologies such as networked backup, LAN-free backup, and serverless backup—and then the evolution stopped. While each of these methods mitigates one or another of the backup challenges, its limitations quickly became apparent with use: each method leaves a set of problems unresolved, deteriorates performance in new ways, relocates inefficiencies, and/or introduces significant additional financial burdens of its own. (These backup methods and their limitations are discussed in the later section “A Lesson In Types of Backup,” starting on page 5.)

The Total Backup Solution

So here's the good news: there are finally **state-of-the-art tape-based and disk-based enterprise backup and recovery technologies** available that completely solve these pervasive challenges: two innovative methods of accelerating tape backup – **HyperTrac Backup** and **Zero-Impact Backup**, and an advanced method of point-in-time disk journaling, called TimeMark, that works in conjunction with replication to deliver two disk-based backup solutions– **Continuous Online Backup** and **Continuous Nearline Backup**.



While enterprises can choose to do tape- or disk-based backup alone, the optimal backup paradigm for most cases now entails using a tape-based solution in concert with a disk-based solution. This provides the most rapid and effective recovery from both “soft” and “hard” errors, in the near term (use disk-based solutions) and in the long term (use tape-based solutions).

This paper explains these new methods and their benefits, and also provides a detailed discussion in backup methods – from the traditional methods of the past to the new methods of the present, with all the pros and cons of each.

For the last several years, although innovation was strong in other areas of storage technology, no major new breakthroughs had been made in the backup/recovery arena—despite much industry hype to the contrary. Tape backup had, and still has, its rightful place in enterprise data protection, but there was an obvious need for a more robust backup paradigm.

One of the most serious shortcomings of all previous backup paradigms had been the perpetual need for administrators to go through the involved, time-consuming, and often unreliable tape restore process every time the all-too-familiar need arose to recover data (typically single files or records) discovered lost to “**soft errors**” that had occurred in the short-term (i.e., within the last hour/day/week)—errors such as accidental deletion, file corruption, viruses, etc. Not to mention all the priceless data that would be lost – much of it irreplaceably – when they would go back to data on tape that was a day old, or older.

It was becoming apparent that with the much faster response time of disk (and no need to fuss with moving tapes around), it should be much easier and faster to restore such lost files or volumes directly from disk rather than tape. Since soft errors account for *more than half* of all data lost (*see pie chart on the previous page*), restoring from disk would clearly save huge amounts of time for both administrators and the employees waiting on the sidelines for their data to be restored, with the happy side effect of greater enterprise productivity. But back in the day, disk was still prohibitively expensive and the sophisticated software necessary to make disk-based backup and recovery practical, granular, fast, low impact, and cost-effective wasn't available...yet.

Now Available: Advanced Disk-Based Backup Solutions for Immediate Recovery...

Today there is a great alternative to tape recovery for the purpose of immediately restoring data lost to soft errors in the short term: disk-based backup. Disk itself has become cheap, its price per MB continuing to fall drastically. And the solution for practical, cost-effective facilitation of granular disk-based backup has finally arrived: **TimeMark**, the continuous snapshot option of FalconStor's IPStor storage networking infrastructure software suite. It is the driving force behind the most advanced disk-based backup available.

TimeMark enables administrators to create multiple consecutive delta-snapshots (“TimeMarks”) of live data and stores them on disk. The ideal complement to traditional tape backup, the point-in-time disk journaling that TimeMark performs minimizes the risk of losing data in between tape backups.

TimeMark is the foundation for FalconStor's two extremely effective and easy-to-use solutions for **no-impact, anytime backup of live data to disk – Continuous Online Backup** (local disk-based backup) and **Continuous Nearline Backup** (remote disk-based backup) – both of which deliver instantaneous backup and **immediate recovery of mission-critical data**. (*See the section “Disk-based Continuous Backup: Enhancing Tape Backup and Beyond” on page 13 for more details.*)

...And Accelerated Zero-Impact Tape Backup Solutions for Long-Term Recovery

Most enterprises perform tape backup, and will continue doing so for the foreseeable future. After all, tape plays an important role in data archiving for legal and other purposes, data warehousing, transporting data to a new location, and/or as a second line of defense against hard errors in addition to methods of disk-based backup, mirroring, and/or replication. And while disk-based methods are superior for near-term recovery, disk capacity is ultimately limited— a given disk scenario can provide about a week's worth of protection (for example, you could efficiently store a week's worth of hourly TimeMarks). On the other hand you can easily create and store an infinite amount of tapes.

The widespread use of tape makes it crucial to find ways to enhance the performance of tape backups to relieve the outrageous burdens placed on administrators (e.g. doing an 8 hour backup in 5 hours) or restrictions placed on companies (e.g. the inability to increase uptime because of the need to accommodate backups). To this end, FalconStor has developed two revolutionary backup consolidation and acceleration

solutions, **HyperTrac Backup** and **Zero-Impact Backup**, both of which **eliminate the backup window by permitting full, incremental, or differential tape backup of live data anytime, with little or no impact on the application servers or the LAN.** (See the section “HyperTrac Backup and Zero-Impact Backup: The Revolution!” on page 10 for more details.)

Both methods offload backup services from the application servers and the LAN for truly windowless backup. The main difference between the two methods is that while HyperTrac Backup dramatically accelerates enterprises’ *existing file-by-file or block-level backup software/backup server scenario*, Zero-Impact Backup accelerates block-level backup using *certified block-level backup software that is hosted on the IPStor Server*, centralizing backup management there. As explained later, the TimeMark Option (which is the foundation for disk-based backup as mentioned before) also plays a pivotal role in bringing tape backup to unprecedented levels of performance and granularity, by providing the means to perform continuous snapshot operations and track changes at the block level.

Snapshot Agents Ensure Transactional Integrity and Point-in-Time Consistency During Backup

Transactional integrity and point-in-time consistency of databases and messaging data stores are vital elements if backups are to actually be viable in the event of a restore situation. If the backup does not contain transactionally accurate database records, or does not reflect an accurate portrait of a message store, the results could be very harmful to business. Both **transactional integrity and point-in-time consistency are ensured during all backup processes** with IPStor’s application-aware **Snapshot Agents** (currently available for Microsoft® Exchange, Oracle®, Lotus® Notes, Sybase®, IBM® DB2 UDB, and SQL Server). The Snapshot Agents also preclude the necessity of putting databases in backup mode for extended periods of time – quiesce time can now be measured in seconds rather than hours.

The Freedom to Choose the Best Backup Scenario for You

HyperTrac Backup and **Zero-Impact Backup**, **Continuous Online Backup** and **Continuous Nearline Backup**. The unique features of each of these four backup methods adapt themselves easily to the different storage environments and requirements of diverse enterprises.

- Easily integrated with each other, these four revolutionary backup solutions solve the full gamut of challenges while lowering the total cost of owning a backup solution.
- There’s no more worrying about backup windows or lengthy recovery times.
- Gone are the high costs and excessive management overhead associated with push agents, open file agents, and backup software database agents.
- Utilizing a highly advanced snapshot-based mechanism and backup paths that entirely avoid impacting application servers, all four of these state-of-the-art methods deliver super-fast, anytime backup of live data with transactional integrity and point-in-time consistency.
- When disk-based backup is used in between tape backups, the risk window of data loss is dramatically minimized.

With effective disk-based and tape-based backup solutions to choose from for fully **customizable total data protection**, administrators now have the ability to establish both a long-term and short-term strategy for their enterprises and can *finally stop losing sleep* over their backup issues.

A Lesson in Types of Backup

A lot of bad days can be avoided by knowing what kinds of backup/recovery methods are available. So let's take some time now to look at the old way of backup and move up the evolutionary ladder to more recent methodologies and groundbreaking backup methods.

Backup has come a long way since the early days of the IT industry – direct-connect backups for each app server are now not the only option. Each successive backup strategy developed to address shortcomings of previous strategies in either backup speed/performance or impact on app servers. Today, there are several major types of backup that most data center managers have implemented or are implementing:

Tape-based backup

- **Dedicated backup network** – backup occurs across a network to a dedicated backup server to tape, with each application server facilitating movement of its data to the backup server (p.7)
- **LAN-free backup** – backup occurs across the SAN with app servers sharing library resources (p.7-8)
- **Serverless backup** – with backup software on each applications server, backup occurs across the SAN with the aid of a proprietary data mover that utilizes the extended copy command to move data from storage to tape. (p. 9)
- **HyperTrac Backup** – using existing backup software on an existing dedicated backup server, file-by-file or block level backup of live data to tape occurs directly between the storage device and the tape library resources at speeds up to five times faster than standard backups, with no effect on the LAN or on the application servers (p.10)
- **Zero-Impact Backup** – with the IPStor Server hosting the backup software for centralizes backup management, block level backup of live data to tape occurs directly between the storage device and the tape library resources at speeds up to five times faster than standard backups, with no effect on the LAN or on the application servers (p.11-12)

Disk-based backup

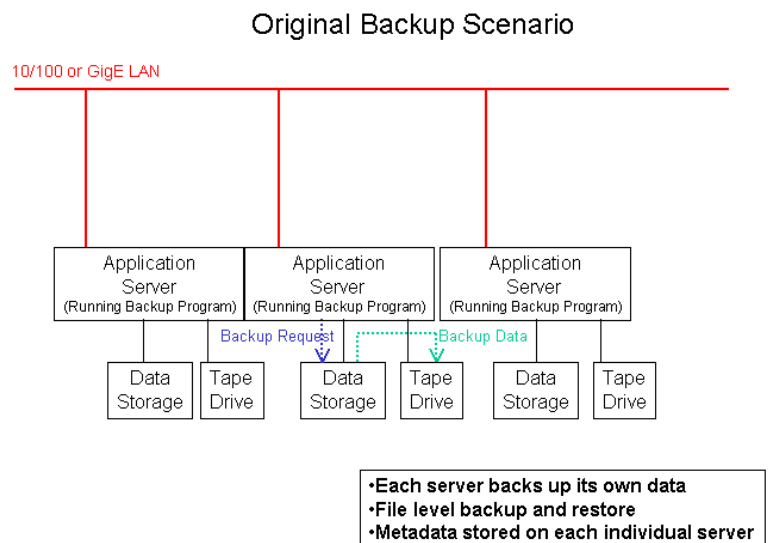
- **Continuous Online Backup** – data is replicated locally, and then delta-snapshot-based backup of live data occurs directly from storage to disk with no impact on the LAN or on application servers (p.13)
- **Continuous Nearline Backup** – data is replicated to a remote site, and then delta-snapshot-based backup of data occurs directly from storage to disk at the remote location (p.13-14)

IPStor software from FalconStor enhances the dedicated backup and LAN-free backup scenarios. More importantly, it is the enabler for the newest and most effective of backup strategies, namely tape-based HyperTrac Backup and Zero-Impact Backup (both of which far outperform serverless backup), and the two kinds of disk-based Continuous Backup.

Back in the Day: Traditional Backups

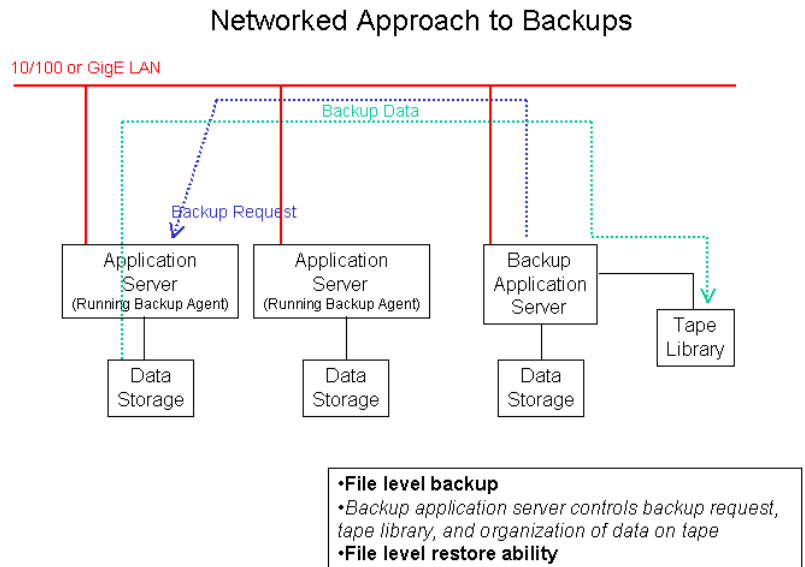
The dawn of backup

Once data became a mission critical resource for companies, the need to store it and to protect it became apparent. Initially, backups entailed each app server having its own dedicated backup tape drive. While this allowed for fast backups of each server, it also created a management nightmare. Each tape device had to be administered individually, removing and storing tapes done by hand, and each individual server had to have its own copy of the backup app. As organizations grew and the amount of data and amount of servers to be backed up increased, the costs of administration hours, purchasing backup applications at every server, and the likelihood of human error became obstacles to ensuring data integrity.



Networked backup methods emerge

To deal with the drawbacks of traditional “single server, single backup” scenarios, data center managers began to adopt a networked approach to backup. In this scenario, all the application servers are connected via a local area network (LAN), with a single server (usually dedicated) designated as the backup application server. Backup software is loaded on the backup application server, which initiates and controls all file backups. The backup application server sends backup commands across the LAN to the server to be backed up. Files from that server are sent back across the network to the backup server, which then sends them to the tape library.



In this setup, backup resources (robotics and tape drives) and metadata (organization of the data on tape) are controlled by the backup application software installed on the dedicated backup server. The backup application communicates with backup agents that are installed and running on the other application servers on the network. These agents allow the backup server to facilitate backup schedules, access directory structures, and initiate movement of data, enabling both backup and restore at the file level.

Using this scenario, most organizations implemented a weekly full backup (complete system backup) and daily incremental backups (changes since the last full). Backups were scheduled during business off-hours in what was called the “backup window,” a time where there was very little network traffic from normal business usage. However, if any files or databases happened to be in-use, these had to be skipped during the backup.

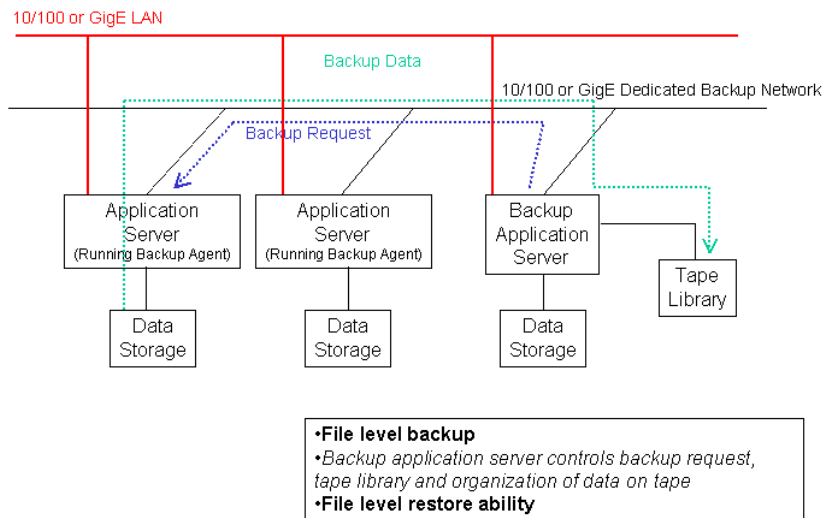
As the amount of data to be backed up increased, and the business pressures of having to operate on a 24x7 basis became the norm, the so-called “backup window” shrank to almost nil. Problems arose when **backup traffic extended beyond the backup window and began clogging the LAN, taking bandwidth that was needed for normal business traffic.** Worse still, **active files, whether they were small Excel files or the entire Exchange data store, were skipped**—flouting the whole point of having a plan for critical data back up.

The Next Few Generations of Backup

Dedicated Backup Network

To address the bandwidth issue, today many traditional backups use a dedicated backup network (in some cases GigE but usually 10/100) to separate backup traffic from business traffic. The dedicated backup server still hosts the backup application and a backup agent still resides on each application server that it is attached to, only now all metadata and file traffic occurs over the dedicated backup network.

Traditional Backups with Dedicated Backup Network



Use of a dedicated backup network eliminates backup traffic from the normal business LAN. It also allows the consolidation of backup resources and backup management. However, it does not address the **slow data transfer speeds** of existing network connections or the expense of the backup software package and its agents. Moreover, **this backup scenario also actually increases the processor load on each application server**. During backup, each individual server has to facilitate the movement of its data to the backup server. This is in addition to the **inability to backup active files** and the **necessity to quiesce databases for very long periods of time in order to maintain transactional integrity** for the duration of the backup.

LAN-free Backup: An Evolution

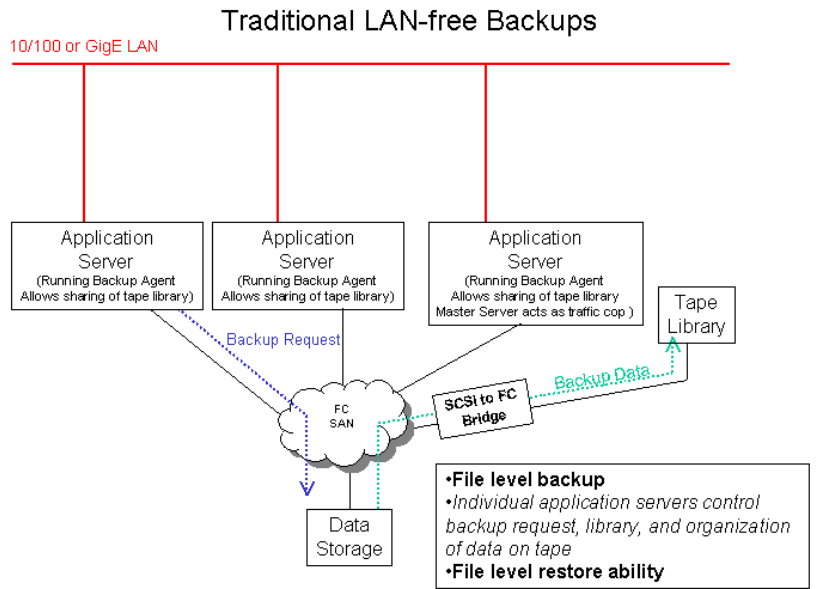
With the introduction of the storage area network (SAN) came the idea of LAN-free backup capability. LAN-free setups enable backups to occur *without* a dedicated backup network *and without negatively* impacting the existing business network's bandwidth. In this scenario, the tape resources are attached directly to the storage network through a switch (most libraries now support Fibre Channel connectivity directly or through a built-in or external FC to SCSI bridge). A SAN-enabled backup application allows each application server to share the tape library through its SAN connection without use of a dedicated backup server.

In LAN-free backups, each application server runs a copy of the backup software with one server acting as the main database for the metadata about the tapes. Each application server is attached to the library resources via the SAN and each shares ownership and control equally. An individual application server initiates backups by sending command messages to the robotics and tape drives as needed. The tape drives and robotics handle potential device conflicts by issuing SCSI reserve/release commands when devices are in use, effectively blocking access to busy devices. Metadata is sent to a master database located on a

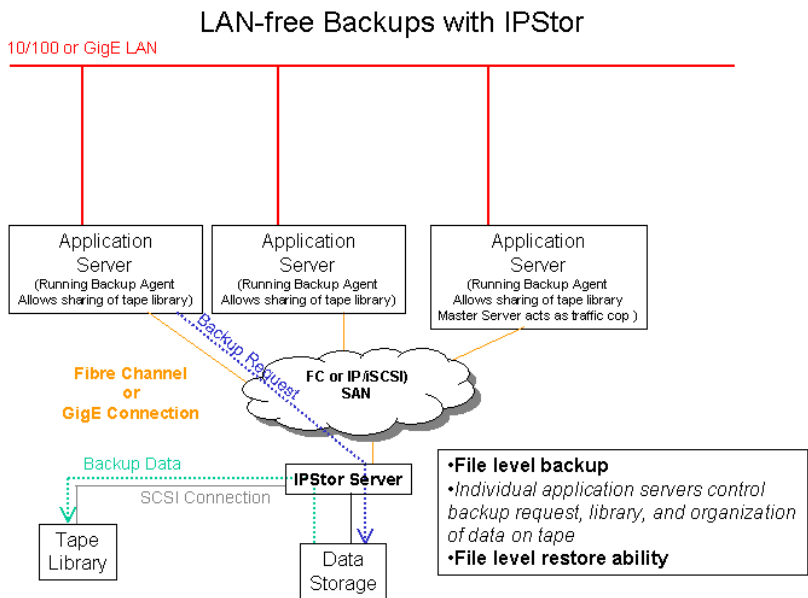
master server. This metadata is available to all servers through the backup agent running on them and allows for a central view of all resources.

Backup commands are issued from individual application servers, but data transfers occur directly between the storage device and the tape library. No command or data traffic is sent across the LAN, hence the term "LAN-free" backup.

LAN-free backups address the issues of data transfer performance because transfers occur across the high speed SAN instead of traditional LAN networks. The data is now transferred across a 100GB/sec Fibre Channel or Gigabit Ethernet SAN instead of a 10/100Mbit IP LAN. It also eliminates backup network communication across regular LAN channels. These two advantages solve several of the performance and LAN traffic issues that have plagued IT backup strategy for some time.



What LAN-free backups do not address is the **high cost of purchasing and maintaining the backup application plus its agents for each application server**, as well as the **additional processor load required for moving the data between the disks and the tape resource**. Still, by comparison to the traditional way of doing backups, many enterprises have decided that the advantages that this scenario brings far outweigh its associated costs.



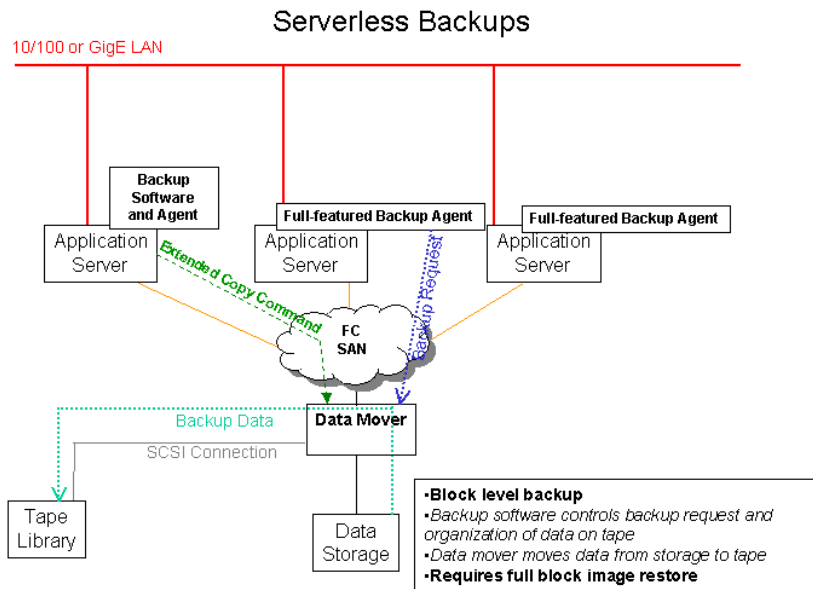
If an enterprise chooses this backup model, IPStor fits well into the LAN-free environment. The application servers still control all backups and data is still transferred directly from disk to tape. The difference is that the data now resides on IPStor-managed drives. One advantage to using IPStor in this scenario is the cost savings from using a SCSI-attached tape library on the IPStor Server instead of a Fibre Channel bridge.

Serverless Backup: the Evolution Continued

Some backup software vendors have moved toward support of what is known as the serverless backup model. This model makes use of what is known as the “extended copy command” along with a “data mover” to move data with supposedly minimal impact on the application servers. Rather than having the application servers “push” the data – which creates substantial performance impact on both CPU and memory – the backup server compiles a list of blocks that need to be moved to tape. This list of blocks is passed to the data mover, which does the heavy lifting of transferring the data from disk to tape.

For serverless backup to work, a separate costly proprietary device is required to sit on the SAN and function as the data mover, which the backup applications use as the facilitator for the extended copy. More often than not, the company that makes such a device (usually a Fibre Channel router or bridge with extra capabilities) is not the same company that makes the backup software, and unfortunately the level of compatibility between the various brands of backup software and the various data movers is low. This is in part because hardware vendors have different implementations of the extended copy command. If a software/data mover match, albeit expensive, is found, the impact to application servers is lessened as data is moved directly from disk to tape. Data traffic is effectively eliminated from the LAN and SAN, and the processing cycle imposed on the application servers is reduced.

The constraints associated with traditional backups are somewhat decreased. The backup software at each application server issues the command to the data mover and then removes itself from the data path. In this way, the data mover and the source and destination devices are involved in handling the data traffic. But what seems like low-impact at first glance really isn't...



The problem with serverless backup is that the **application servers are still required to run full-fledged, intelligent components of the backup software** that submit the backup request as well as **obtain periodic updates on progress**. These factors have a **significant impact on the CPU and memory of the application servers, so the backup is not really “serverless.”** The backup agents are also **expensive and must be managed by an administrator**.

Performing HyperTrac Backup or Zero-Impact Backup using IPStor circumvents all of these issues.

HyperTrac Backup and Zero-Impact Backup: The Revolution!

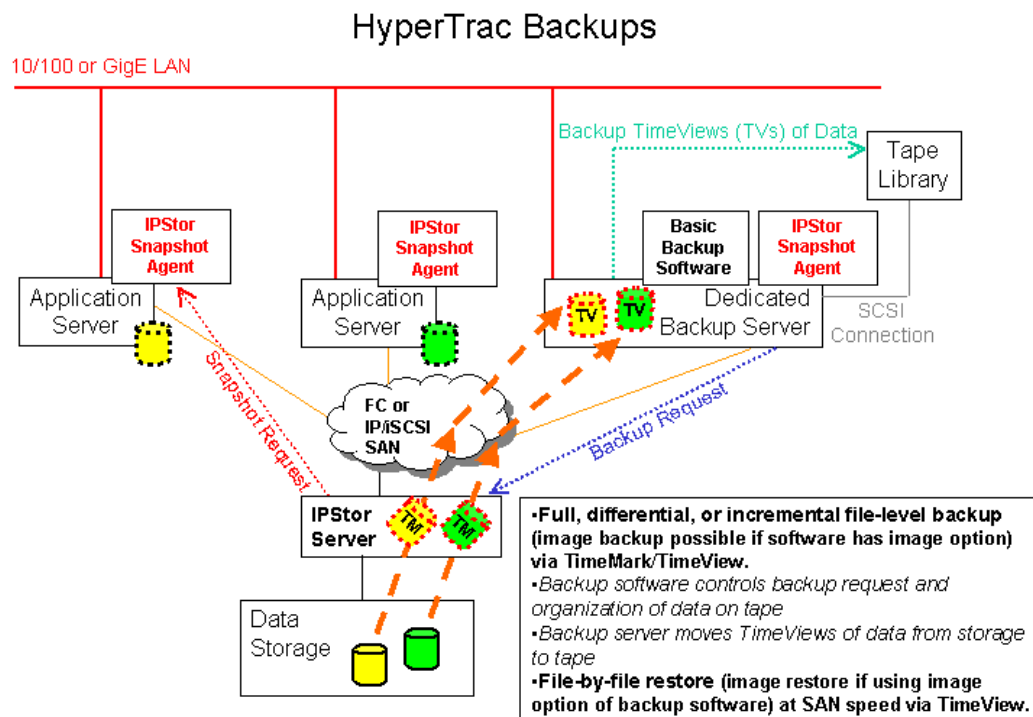
With FalconStor's HyperTrac Backup and Zero-Impact Backup solutions, which employ IPStor to turbo-charge third party backup software, enterprises can achieve unprecedented results for their tape backups:

1. **Perform both file-by-file and block-level backups *without impact to the application servers and LAN.*** Before IPStor, both types of tape backups always heavily impacted the app servers and the LAN.
2. IPStor makes it possible to **perform *incremental or differential block-level backups, in addition to the usual full block-level backups***; before IPStor, it was only possible to perform *full* block-level backups. And there is no impact to the app servers or the LAN whether full, incremental, or differential backup is being performed.

HyperTrac Backup – up to 5x faster backup with no backup window, using existing backup server

For high-speed file-by-file or block-level backup with existing backup software running from an existing backup server

Fast, highly cost-effective, and painless to implement, FalconStor's HyperTrac Backup solution **eliminates the backup window and dramatically accelerates file-by-file backup¹ – making it up to five times faster** – while allowing administrators to continue to **use the existing off-the-shelf backup application they know and like on their existing dedicated backup server**. Backup does not impact the LAN or application servers' memory or CPU for truly windowless copying of data to tape. Restores are granular and very fast.



How does HyperTrac Backup work? IPStor's **HyperTrac Backup Accelerator** works with the **TimeMark Option** to allow an existing dedicated backup server to back up a **TimeView** – a mountable TimeMark snapshot image – of a *volume that can be concurrently in use* by a host server. The HyperTrac Backup

¹ Block-level backup is possible for those who want to do it, by using the image option in backup software that has it.

Accelerator automates the whole process whereby the dedicated backup server backs up/restores each app server's IPStor-managed disks via TimeMark/TimeView at the speed of SAN. Since the backup software backs up the point-in-time snapshots of data (TimeMarks) directly from SAN storage, backup processing is offloaded from application servers and the LAN, eliminating the backup window. The interaction of these two options is what enables IPStor to deliver the administrator's choice of **full, incremental, or differential** high-speed file-by-file backup of live data, anytime.

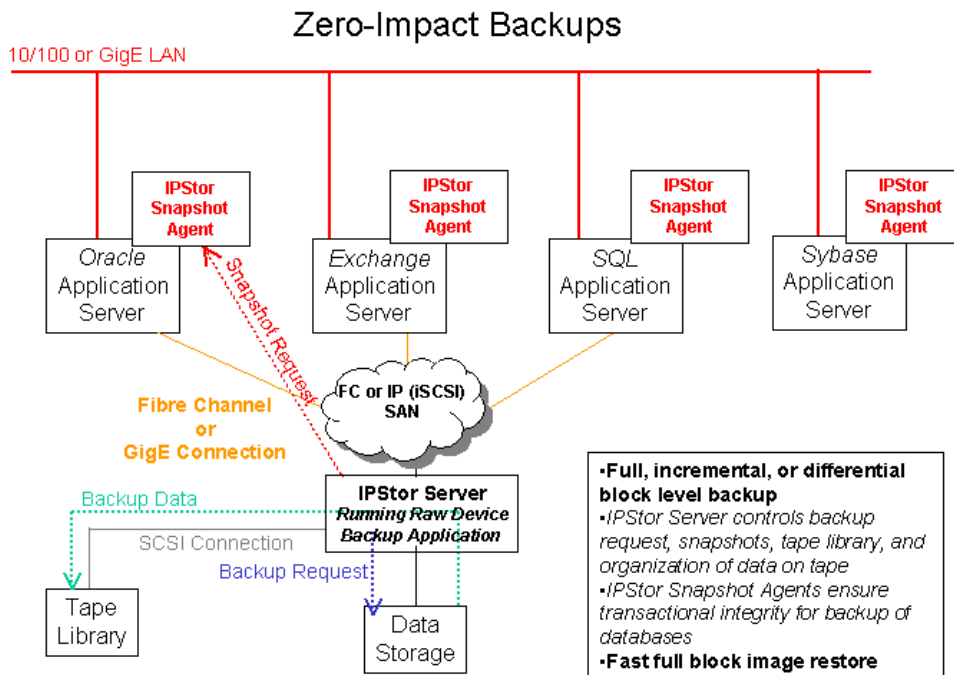
Only one copy of any standard backup software is necessary, running on the dedicated backup server. Since there is no more need for app servers to host expensive push agents, open file agents, tape library sharing options, SAN options, and backup software database agents, **TCO is minimized and backup management** is extremely simplified. All options of the backup software remain functional (including image backup/restore if the backup software has an image option) while additional functionality is granted.

Transactional integrity and point-in-time consistency of databases and messaging data stores is ensured during backup with IPStor's application-aware Snapshot Agents. As HyperTrac Backup is a snapshot-based process, the Snapshot Agents eliminate the necessity of putting databases in backup mode for extended periods of time; **quiesce time can now be measured in seconds**, rather than the hours that traditional methods of tape backup have required.

Zero-Impact Backup – up to 5x faster backup with no backup window, using the IPStor Server

For high-speed block level backup with backup software running on an IPStor Server

Another backup consolidation and acceleration solution offered by FalconStor is Zero-Impact Backup (ZIB). Zero-Impact Backup delivers **full, incremental, or differential block level backup of live data at high speed, with little or no impact on application servers or the LAN, and with unrivaled fast full block level restores**. ZIB allows FalconStor-certified third party backup applications that support raw device backup to run on the IPStor Server for centralized backup management. This solution is highly cost-effective, as only one copy of backup software is necessary. There is also no more need for expensive push agents or open file agents running at the app servers. This further minimizes TCO and offers simplified backup management.



How does ZIB work? Tape backup using IPStor's **Zero-Impact Backup Enabler Option** outperforms serverless backup by completely eliminating the need for application server intervention during backup and restore operations. There is no added processor or memory load on the individual application servers and there is no need for additional backup application agents, since both command and data traffic flows happen without any application server intervention whatsoever (i.e., control, metadata, and data flow happen without any application server activity). The enabler for this scenario resides in the use of an intelligent, cost-effective data mover to handle the tape resources, the metadata, and the movement of data to tape.

That data mover is the IPStor Server, running an off-the-shelf raw-device-enabled backup application (such as those from Veritas, Legato, and BakBone). The IPStor Server presides over control of the backups, control of the tape library, and metadata organization on the backup tapes. There is no interaction required with any app servers during a backup sequence. During a backup, the IPStor Server grabs a full image, block level copy of the requested data and moves it directly to the tape library. No app server requests occur and no information is exchanged between the data mover and the app servers. Once initiated, the data copy happens directly from the storage to the library resource. App servers on the SAN benefit from performance increases and the elimination of overhead associated with backup/restore operations because the command and data paths are rendered exclusively local to the IPStor Server. This results in the most optimal data transfer between the disks and the tape, and achieves net transfer rates that are limited only by the disk's or tape's engine. The backup process automatically leverages IPStor's snapshot engine and application-aware Snapshot Agents to guarantee point-in-time consistency and transactional integrity.

Because no directory structure copy is needed, data is written at a block-by-block level. This enables lightning-fast backups. There is very little overhead and the backup device can be optimized to block size performance. Using IPStor as the data mover, tape device speeds for backups are possible. In scenarios where extremely fast data backups are required (as most scenarios are quick becoming) there simply is no better solution.

Data center downtime carries an hourly price tag of anywhere from \$90,000 for a catalog retailer to \$6.5 million for an online brokerage.

(Data Base, Inc.)

Another huge advantage is the *cost savings* possible through eliminating the need for agent software on each application server. Because the data mover does not require any information from them, the application servers do not need to have any backup agents installed. In some scenarios this can provide a significant cost savings to the customer. Zero-impact backups also eliminate any additional processor load on the application server. All data movements are controlled by the data mover with no impact on the application servers, hence the term "zero-impact" backup.

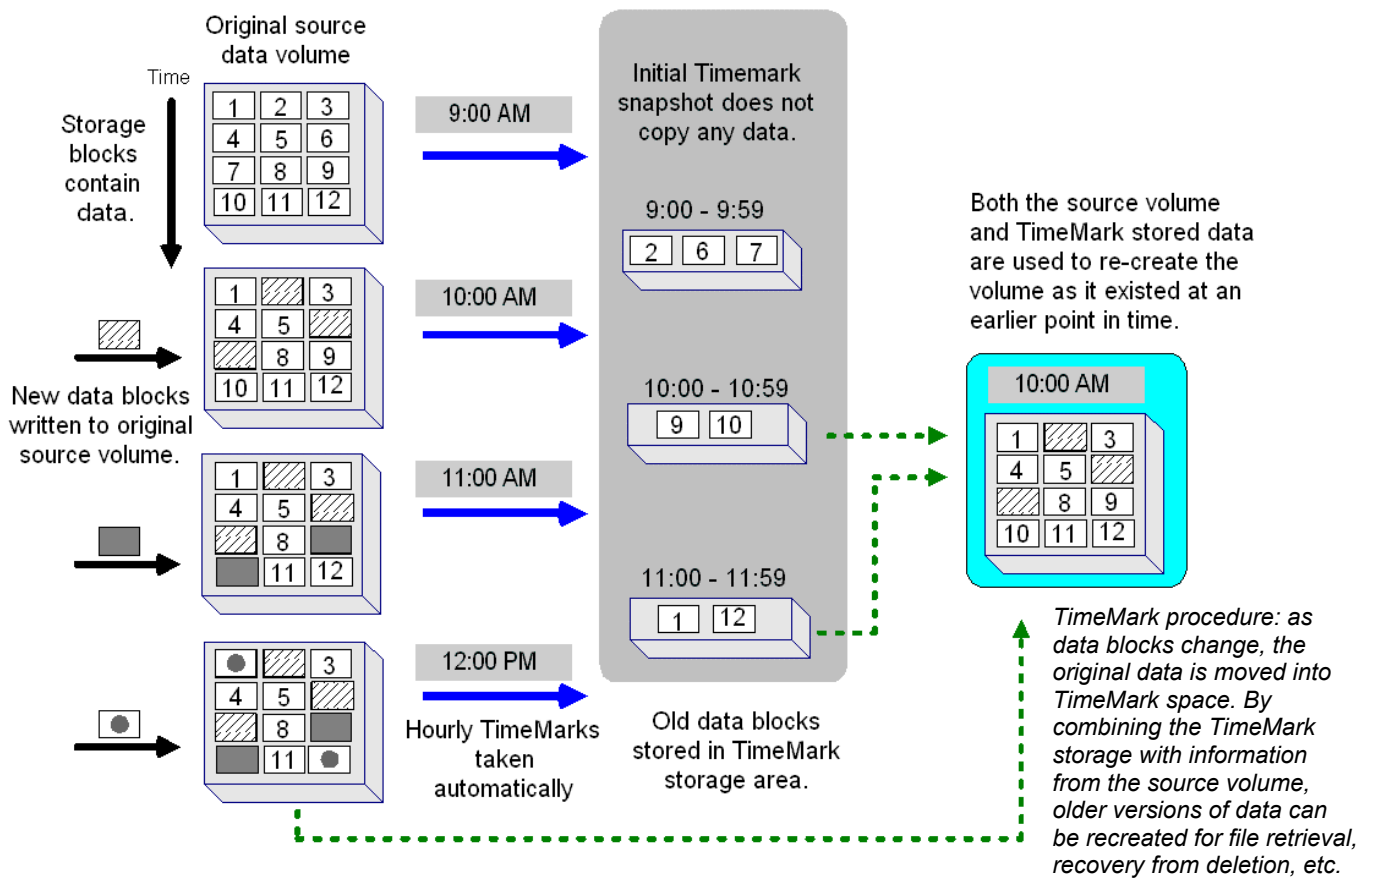
Because zero-impact backup is a full image, block level backup, the restore is also a full image block level restore. Therefore, the speed of restoring an entire volume is also at the fastest speed. This is where IPStor's revolutionary TimeMark/TimeView feature comes in, allowing for near-line file-level backup and restores in the shorter term, without impacting the LAN or application servers, and without the headaches of searching on tape for that file that someone just accidentally deleted.

Disk-based Continuous Backup: Enhancing Tape Backup and Beyond

Continuous Online Backup and Continuous Nearline Backup

TimeMark is the foundation for **Continuous Online Backup** and **Continuous Nearline Backup**. Both methods deliver full, incremental, or differential automated instant backup to low-cost disk, without involving backup software. There is no impact to application servers, making anytime backups possible—the issue of the backup window is altogether eliminated. Both protect from “soft errors” and deliver immediate, easy recovery of individual files/records or volumes. The difference is that **while in the Continuous Online Backup scenario the primary storage and TimeMarks are local, in the Continuous Nearline Backup scenario the TimeMarks are taken off the replica data at a remote site (centralized data center or disaster recovery site), providing an extra layer of protection.**

IPStor's TimeMark feature does NOT replace the need to have tape products. Disk-based Continuous Backups complement tape backups and help to reduce the need to perform incremental tape backups and minimizing the risk of data loss in between tape backups. In fact, FalconStor actively works with leading tape system hardware and software vendors to collaborate on how IPStor's technology can be further leveraged to enhance and simplify tape backup methods.



What is the purpose of TimeMark? Near-instantaneous recovery of single files—or an entire volume—back to a known-good point in time. How does Continuous Backup work? TimeMark enables IPStor to create periodic, scheduled point-in-time copies of data volumes to achieve online journaling of disks. TimeMark volumes contain only data changes (what FalconStor calls “delta-snapshots”) and therefore do not take up a

significant amount of storage. You can schedule a TimeMark every day or even every hour, and be able to have hourly recoverability of lost files or the entire disk. Up to 256 TimeMark copies can be maintained per volume.

Using IPStor's TimeView feature, an extension of the TimeMark option, administrators are provided with the tools to freely create multiple and instantaneous disk images of an active data set or replica (created using IPStor's Replication option). The data set and/or replica copies can then be assigned to an application server with read/write access for independent processing, all while the original data set is still actively being accessed/updated by the primary application server. TimeView is an extremely valuable tool since it enables users to create a disk image out of any TimeMark, making the contents as of that moment of time fully accessible (read and write) for recovering lost files or for data processing.

A small amount of storage (not necessarily from the same storage as the primary disk) is pre-defined to hold the incremental changes to the disk (block level, not file level). When a TimeMark is created (automatically per the schedule, or on demand using IPStor's Java Console or command line interface) the pre-change data (original contents) are held in this area. Therefore, it is not necessary to have 100% of the amount of disk space in order to create a TimeMark. The typical amount is around 1% to 5%, though this depends on the amount of activity (# of new writes) to the disk.

Instant, and more granular recovery

IT managers normally have to worry about two types of data loss problems: hardware (disk) failure and "soft errors" which include accidental deletions, disk partition or file system corruptions, application level software glitches, or virus attacks. While RAID, mirroring and replication take care of hardware problems, performing effective and frequent backups is the only defense against "soft errors." But how frequently are disks backed up to tape? Typically, there is a block level full backup on the weekend, supplemented with Monday, Tuesday, Wednesday ... incremental backups. For large databases, even incremental backups (file based) are not effective because the data is changing continuously. With TimeMark, you can simply schedule an hourly virtual TimeMark snapshot (5 days per week, 24 TimeMarks per day = 120 TimeMarks per week) and can even automatically recycle the space on a weekly basis.

Suppose that on Thursday at 5:32 PM, you discovered that a massive virus attack has corrupted your entire disk, and a virus scan/repair tool for that particular virus isn't yet available. You trace the problem to an infected program that you ran at 5:02 PM. Using IPStor, you can simply right-click on that infected drive, select "TimeMark -> Rollback" and select the "Thursday 5:00PM" TimeMark. Instantly, the contents of the entire drive are rolled back to the exact contents as of 5:00PM, like a time machine. This effectively reverses the mishap (running the infected program) that occurred at 5:02 PM.

Without IPStor's TimeMark option, the above scenario would require restoring the entire disk from tape. If a file-level based restore is performed from tape using third-party backup software with a restore rate of 20GB/hour, the restore of a 100GB drive would take five hours; the restore of a 500GB drive would take more than one day. Using IPStor's TimeMark option, the restore of the drive can be completed in minutes!

Equally important: when using a tape backup, the data is typically from the previous day. So, an entire day's work might not be available on the tape. With TimeMark set for hourly snapshots, the most you can lose is an hour's work. Greater granularity – even as much as a snapshot per minute – is also possible. Furthermore, a TimeView (a mountable TimeMark) can be instantly assigned to all IPStor-supported application servers just like a disk volume. This enables the administrator to get back the entire contents of the disk as of that particular point in time, thereby allowing recovery of a few files instantly, without the need to duplicate the entire volume.

Conclusion

The price of an effective backup/recovery plan pales in comparison to the cost of lost data and downtime. IPStor offers enterprises the industry's most comprehensive, advanced, cost-effective, easy-to-deploy, and easy-to-use solutions. With IPStor, the umbrella of data protection has finally widened to an extraordinary degree, so that enterprise data can now be protected from every category of threat. IPStor enhances existing backup procedures, integrates seamlessly into existing backup environments, and enables entirely new paradigms for high-speed backup and near-immediate recovery that combine storage consolidation and acceleration with advanced Zero-Impact Backup and disk-based Continuous Backup methods.

With IPStor you get the granular, immediate, on-line recovery needed to get back on track right away after accidental deletions, hacker attacks, or viruses, or any of the innumerable things that might go wrong; you get lightning fast backup that doesn't impact your application servers; and you gain the unprecedented freedom to choose to perform full, incremental and/or differential file-by-file or block level tape backup that can occur unnoticed behind the scenes even as your business zooms along.

The strategies and solutions to fully protect enterprise data exist at last; all that's left is to implement them. The old guard in backup has made way for the new. Whereas there have been no significant improvements in backup software technology for years, IPStor takes existing off-the-shelf backup software and revs it up for unprecedented high-speed, no-impact, granular backup. IPStor also integrates disk-based and tape-based backup, making it easy for enterprises to implement a backup plan that successfully, efficiently, and cost-effectively protects against data loss from both soft and hard errors, in the short term and the long-term.