

---

---

## LiveVault® Software Planning and Best Practices Guide

---

---

### 1. Introduction

The LiveVault Software™ gives you the ability to implement an entirely disk-based backup and off site archiving solution custom tailored to the needs of your business. This guide will help you plan your LiveVault deployment and help you recognize some of the best practices that reduce risk, provide greater reliability, and provide your desired recovery time experience.

In outline, a LiveVault Software deployment involves:

- Understanding the replication strategies
- Determining the vault configurations, the number of vaults and their locations
- Estimating backup and recovery times and assessing your risks
- Managing vaults, vault loading and bandwidth

### 2. Vaults and Replication Strategies

#### **Vaults**

A LiveVault vault is a Windows 2003 server with up to 6TB of disk space<sup>1</sup>. You have great latitude in selecting the configuration. For example, you can use inexpensive RAID 5 SATA disk enclosures, a typical choice, or SAN-based storage, or many other storage configurations<sup>2</sup>.

The amount of vault space you need depends on the amount of data to be protected, and the length of time you want to retain backup versions<sup>3</sup>. With the LiveVault Software you configure the history retention options that are right for you. For example:

- For 30 day history, plan on 3 times the amount of protected data
- For 1 year history, plan on 5 times the amount of protected data
- For 7 year history, plan on 7 times the amount of protected data

If you need more than 6TB of vault disk space, you will have multiple primary vaults. The LiveVault scalability strategy is to divide the backup load between smaller (inexpensive) vault servers rather than to rely on a single large server. This reduces the impact of a vault failure and allows you to expand your LiveVault deployment gradually as your needs grow. It also allows you to purchase vaulting capacity using “commodity” servers rather than requiring high-end systems.

In a LiveVault deployment there are no tapes. This eliminates the primary cause of unreliable backups, dramatically reduces the labor and on-going operating costs, and provides much better reliability for longer term archiving. What confidence can you have that a tape written today (with today’s software) will still be readable seven years from now?

Because there are no tapes<sup>4</sup>, you should plan to have at least two copies of all vaulted data for high availability. Normally there will be at least a primary and secondary vault. A primary vault is one to which

---

<sup>1</sup> Full specifications are available separately.

<sup>2</sup> If you want to use network attached storage, such as iSCSI or mapped drives, please discuss this with your Iron Mountain Sales Engineer.

<sup>3</sup> The vaults need a significant amount of free space as temporary work space for their operation. This requirement is reflected in the 3x, 5x and 7x factors. Only 60% or less of the disk space should be used for storage of backup data. The actual amount of storage you need may differ from the examples given.

agents (or TurboRestore appliances) send data. A secondary vault receives completed backups from other vaults.

Using just one vault may be viable if you can tolerate the consequences of a vault failure. Vault failures divide into two types: outages, in which the vault storage is not lost, and catastrophic failures where the vault's disk storage is lost. A single vault implementation always involves the use of one or more TurboRestore appliances as well as the vault. In this manner, backups to the appliance can continue even if the vault is down, and in the case of a catastrophic vault failure at least some backup versions remain on the appliance. However, as described later (under "lossy" replication), after a catastrophic vault failure the appliance will only send the latest backup version to the new vault; it will not send all the versions that it has retained.

**TurboRestore Appliances.** As noted above, you can have TurboRestore appliances at some or all of your sites. As with vaults, you provide hardware that meets the specification and Iron Mountain provides the software. With an appliance, backups can be made first to the appliance and the appliance will then backup to the vault, perhaps on a less frequent schedule than backups are done to the appliance. An appliance can keep backup history with the same retention as the vault, space permitting. With normal loading and configuration an appliance will keep 2 to 3 weeks of historical backups for the servers it is protecting. A lightly loaded appliance may have room to keep more history. An overloaded appliance will take itself out of service if it does not have space to keep at least 2 days of backup history. (You will receive a warning email if the capacity falls below one week, and an email alert if the appliance has to take itself out of service.)

For each server that is backing up to an appliance, there is no guarantee that the appliance will have the same backup versions as the vault:

- The appliance will have more versions than the vault when the appliance backups to the vault less frequently than the agents back up to the appliance, and after recovery from a catastrophic vault failure when there is only a single vault.
- An appliance will have fewer versions if there isn't room to keep the oldest versions on the appliance, and after an appliance fails (losing its disk space) and is repaired or replaced. In this case the agents will perform a new complete backup to the empty appliance, but a vault never "down loads" backup versions to an appliance. This is an important point: replication is in one direction only – from appliance to vault.

**Vault-To-Vault Replication ("loss-less" replication).** Vault-to-vault replication is automatic and can occur in either direction. As soon as an agent or a TurboRestore appliance has completed a backup to a primary vault, that vault begins the replication to the secondary vault. Mechanically you implement the vault-to-vault replication strategy by creating *route definitions*. A route says, for example, that vault A should replicate to vault B. There can be more than two vaults in a route, and the same vault can participate in multiple routes.

Vault-to-Vault replication is able to proceed at a rate of 30+GB per hour, bandwidth permitting. The data that is replicated is already compressed and encrypted. So, a vault with 3TB of used disk space requires an estimated 4 days to fully synchronize with another vault, given adequate bandwidth. (This would only be required to establish a new vault, as would happen after a catastrophic vault failure.)

In the event that a vault fails and you know that it will take awhile to get it back in service, through the LiveVault user interface the vault is declared to be out of service. Any agents and appliances that use the vault as a primary vault will be automatically redirected to back up to the next vault in the route as their (new) primary vault.

When the failed vault is back on line, the LiveVault interface is used to place it back in service. The other vaults in the route (either before or after it) will replicate to it any backups that it missed while it was down.

---

<sup>4</sup> Backing up a LiveVault vault to tape is problematic and not supported.

You can manually manipulate route definitions as necessary. For example, if a vault is becoming over loaded, you can perform load balancing by adding another vault to the route used by some of the agents and/or appliances, let the replication to the new vault complete, then redirect those agents and appliances to use the new vault as their primary vault.

**Handling Vault Problems.** The objective of vault-to-vault replication is for both vaults to have all the same backup data (for the routes in which they participate). When a new backup has completed for an agent's backup policy, the primary vault will immediately send the data (the delta backup changes) to the secondary vault. If it can't, because the secondary vault is down or the connection between the vaults is down, it will do the replication automatically when the other vault and/or the connection is restored. If a vault or connection is down for a long time, when the infrastructure is restored there can be a large backlog of data to be replicated. This is called "loss-less" replication. Replication is always done oldest first, newest last.

A problematic scenario is one in which the bandwidth between the two vaults is almost entirely consumed keeping up with the on-going, day-to-day replication when both vaults are operating normally. In this scenario, it's possible that after a vault failure is repaired, the two vaults will never be able to re-synchronize. One possible method of handling this situation is to plan to physically place the two vaults on the same LAN after a failure and then re-separate them after they are again in sync. See the discussion of bandwidth planning later.

During the time that a vault or connection is down, there is a potential risk of losing backups. If a vault fails and its disk storage is lost and then the disk storage on the remaining vault is also lost (a double failure) then you would lose all of the backup history<sup>5</sup>, as well as being unable to make new backups. This is a worst case scenario, but you need to decide what your risk tolerance is as you plan your deployment. There are many options available to you:

- You can have a deployment that allows you to recover from a failed vault as quickly as possible, thus minimizing the time window (and thus the probability) of a problematic double failure. For example, you can have a stand-by vault that you immediately add to the route of any vault that fails so that as quickly as possible you again have two functioning vaults, even while the failure is being repaired or replaced.
- To minimize the impact of bandwidth on vault-to-vault resynchronization times, you can be sure there is adequate low latency bandwidth.
- You can have a third copy of some or all of your backup history, as well as preserve your ability to continue to make new backups, by having a TurboRestore appliance or a third vault. (Typically, a TurboRestore appliance will keep recent backups, but not deep history.)

You do not need to over-design. If you are only planning to keep backup history for a short time, say 14 to 30 days, the loss of backup history may not be a major concern. After such a loss, you would again have complete backup history in 14 to 30 days.

If a connection goes down or a vault fails without losing its permanent storage, then the consequences are often less severe. When the vault and/or the connection is repaired there will be a backlog of recent backups that may need to be replicated, but the entire contents of the vault storage will not need to be sent.

If an agent loses its connection to a vault, no new backups can be made. A remedy is to have a TurboRestore appliance at the server's site so that at least local backups can continue.

**Appliance-To-Vault Replication ("lossy" replication).** Obviously, an agent can only send backup data if there is a connection. If the connection is lost while a backup is in progress, the same backup job will automatically resume when the connection is restored. After this, the next backup will start at its appointed time. (Backups that were missed while the connection was down are simply skipped.)

If the agent is backing up to a TurboRestore appliance, then the appliance can continue to receive backups when the connection is down<sup>6</sup>. When connection is restored, the appliance will only replicate the most

---

<sup>5</sup> Assuming there were only two vaults in the route and no TurboRestore appliances in use.

<sup>6</sup> If a vault goes down, but is not taken out of service, the agent or appliance will not be automatically redirected to a secondary vault. From the agent or appliance view point this situation is just like a dropped network connection. So if a

recent backup version; older backups that completed while the connection was down will continue to reside on the appliance only. (A more common case of this is when the appliance is on a different schedule than the agent. For example, if the agent backs up to the appliance every 15 minutes, but the appliance only backs up to the vault once a night, there will be many backup versions made during the day that only reside on the appliance.) This means that even if a connection is down for a long time, or the connection speed is very slow, the accumulated back log of queued up data is limited to one backup version, and the vault can more quickly obtain a recent backup (the most recent backup), but the vault will not have all the backup versions that were captured on the appliance. This is called “lossy” replication and requires less bandwidth to recover from failures.

**Planning for Off-Site Protection.** A primary consideration in planning a LiveVault Software deployment is to meet reasonable off-site protection goals within an acceptable budget. For some customers the cost of the bandwidth between sites is an issue. With this in mind, here are some options you can consider:

- Can the primary and secondary vaults be at the same location? If the protected servers are at sites remote from the planned location of the vaults, is it an acceptable business risk that in the event the vaulting location is destroyed the backup history for the remote sites would be lost? Another factor is that the remote sites would not be able to make new, current backups until new vaults are brought online, unless TurboRestore appliances are located at the remote sites.
- If some of the servers to be protected are at a planned vaulting site, you will almost certainly need a second vaulting site so that both a server and its backups are not confined to one location.
- If the primary and secondary vaults are geographically separated, can vault recovery still be done locally? If so, the amount of bandwidth needed between the sites will be lower. For example if the two sites are “driving distance” away, then you may plan to set up a repaired or new vault at the same location as a functioning vault, let the vaults re-synchronize, then drive the new vault to its normal (remote) location.
- If the vaults remain geographically separated in a vault failure scenario, can you arrange for a burstable bandwidth contract that allows you to have sufficient bandwidth when needed, but only normally pay for less? (Estimating aggregate bandwidth requirements is discussed in the next section.)
- If you have critical data availability requirements for backup, consider the following options:
  - Have one vaulting location where there are no servers needing protection. Have an appliance at each site with servers to be protected. The primary and secondary vaults are at the one vaulting site. You will have 3 copies of recent backups, but will lose older backups if (and only if) you have a double vault failure within the recovery window after the first vault failure. This is most likely if the vaulting site experiences a site disaster.
  - Have two separate vaulting sites. You have the risk of losing backup history if (and only if) the disk storage of both vaults is lost within the vault recovery window after the first vault failure. The destruction of one vaulting site does not lead to a double failure.
  - Have two separate vaulting sites with a TurboRestore appliance at the server sites.
  - Have three vaulting locations. Two vaults would be at one vaulting site, and one vault at the other vaulting site. You would be more likely to consider this configuration if most of the servers that need protection are located at one of the vaulting sites. Remote server sites might or might not have a TurboRestore appliance. The three vaults insures that you can always recover backup history, even if there is a double failure.

**Typical Site Deployments.** To help make sense of the configuration options discussed above let’s look at typical customer deployments. We will caution you that the flexibility of the LiveVault Software makes it easy to create a very sophisticated but expensive plan. With tape-based backup, most companies only send tapes off site to one location, and are at risk of losing their older backups if that storage site is destroyed, or a particular tape is lost or cannot be read. This level of protection represents standard industry practice, so you

---

vault is down, or will be down, for just a short period of time there is no reason to take it out of service. Agents and appliances will automatically resume when it is back on line.

may want consider whether there are business or regulatory requirements that require meeting a higher standard, even though the LiveVault Software allows you to do so.

*Two Vaulting Sites.* Most customers elect to have two geographically separated vaulting sites. This protects both remote servers (three sites in total) as well as servers located at one or the other of the vaulting sites by providing off site protection for all servers. If some of the servers are at remote sites, then TurboRestore appliances can be used in addition to provide better recovery speeds and to provide local backup even if the connection goes down. (There is no reason to have a TurboRestore appliance at a vaulting site.)

If the vaulting sites are “driving distance” away, many customers plan to use the “co-locate, replicate, then separate” strategy for vault recovery. If not, then having adequate bandwidth for vault recovery is important.

*One Vaulting Site.* If there are no servers that need to be protected at the planned vaulting site, then some customers elect to have both vaults at the one vaulting site, perhaps with a TurboRestore appliance at the remote server sites. This avoids the need to have a second vaulting site, and bandwidth between them.

### **3. Bandwidth**

The bandwidth discussion can be divided into two parts: (a) bandwidth requirements between vaults and (b) bandwidth requirements between agents or appliances and the vaults.

**Bandwidth Basics.** It’s often useful to think in GB/hour rather than Mbps:

- A T1 line (1.5Mbps) moves about .5GB/hr assuming 80% efficiency and no compression
- A T3 line (45 Mbps) moves about 16GB/hr assuming 80% efficiency and no compression

Data sent from an agent to the vault normally compresses at least 2:1 so you would typically expect 1GB per hour on an unloaded T1 line. On the other hand, data sent between vaults is already compressed so the rates above are appropriate estimates.

As an example, if there is a low latency T3 line between two vaults each with 3TB of backup data, and one vault is destroyed and replaced, then replicating 3TB will take about 8 days, assuming there is no contention for the link. (If the two vaults are on the same 100Mbps LAN, then the speed of the vaults themselves will be the determining factor in how fast vaults can replicate. Typical rates are in excess of 30GB per hour.) During the time that vaults are replicating following a vault replacement, other backups will be accumulating on one or both of the affected vaults. So the full time required to achieve re-synchronization depends on how long the replication takes, and on the rate of on-going data change at the protected servers.

With fast links, such as a T3 line, the bandwidth latency time often becomes a significant factor that limits the effective throughput of the link. Bandwidth latency can be hidden by using a large network buffer, if the error rate on the line is very low. The LiveVault software uses the largest possible TCP/IP buffer size (called the “window size”) of 65K. Nevertheless as you can see in the table below, even modest latency times can limit the effectiveness of a high speed link. The table is based on a 65K window size. For example, if the bandwidth latency is 50ms each connection will only be able to move 4 to 5 GB/hr on a T3 line.

When vaults are geographically separated, it is often the bandwidth latency that limits the rate of data movement. On a LAN there is essentially no latency.

## Maximum Transmission Rates with a 65K Window Size for Various Latency Times

RTT in ms	GB per Hour	Equiv. Mbps (approx)
3	71.8	107.7
10	21.5	32.3
15	14.4	21.5
20	10.8	16.2
30	7.2	10.8
40	5.4	8.1
50	4.3	6.5
100	2.2	3.2
150	1.4	2.2
200	1.1	1.6
250	0.9	1.3
300	0.7	1.1
350	0.6	0.9
400	0.5	0.8

**If Possible, Test Your Bandwidth.** Because there are so many factors that can affect actual performance, if you have sites that already have bandwidth installed, you are strongly encouraged to measure the actual throughput by copying or moving a few large files totaling 1 to 5 GB and see what the GB/hr rate really is.

**Bandwidth Between Vaults.** As discussed earlier, vaults replicate one to the other as defined by the routes. All data on the vaults is compressed and so no compression occurs during transmission.

How much bandwidth is adequate between vaults depends on two factors: the bandwidth needed to replicate on-going backups as they are created, and the bandwidth needed to achieve vault re-synchronization following a catastrophic vault failure. A 100Mbps or better LAN always provides adequate bandwidth for both purposes.

*On-Going Backup.* In estimating the bandwidth requirement for on going backup activity, you should estimate the average daily change rate, the expected compression, and the number of hours per day that should be used. An example calculation is shown in the table below:

500 Total GB of original data (uncompressed) that is being protected to the vault  
 10% Estimate for the average daily change rate  
 50 GB expected to change per day, pre compression  
 2 :1 compression assumption  
 25.0 GB per day compressed change to send between the vaults  
 12 Hours per day allowed. For continuous backup, you would use the length of the work day. For nightly backup, use the number of hours allowed each night.  
 2.1 GB per hour is necessary transfer rate.

3.0 Mbps needed to achieve a transfer rate of 1GB per hour of compressed data.  
 (Do not use a different value unless you have a special reason to do so.)

**6.25 Mbps recommended between vaults for on-going backup activities**

In this example:

- The GB total is the aggregate (non-compressed) data being protected to a primary vault. For example, if there are two agents and one is backing up 40GB of Exchange data, and the other is backing up 100 GB of file data, you have 140 GB total.
- To estimate average daily change rate Iron Mountain typically uses 10% for database data, and 5% for non-database data.
- To estimate compression, Iron Mountain typically uses the industry-standard assumption of 2:1.
- The Hours per day reflects the reality that in most cases the delta backup data that needs to be replicated will not be generated uniformly throughout a 24 hour period. If Continuous Protection is being used, then most of the change will occur during the business day. If scheduled backup is being used then most of the change will be captured at night during an anticipated backup window. In either case, you want replication to keep up with the change and not fall too far behind.

If you have multiple vaults at one site replicating to multiple vaults at another site, then replications go on in parallel over separate connections. So if you have a T3 line whose throughput is limited by bandwidth latency to effectively 6Mbps, then this link should be able to handle up to 3 vault replications in parallel, with each connection getting approximately 6Mbps. This is all bandwidth theory, of course, and your results may vary. However, it illustrates the benefits of LiveVault's approach to scalability in that replications done in parallel can allow full use of the available bandwidth.

If you provide bandwidth with multiple links, such as two T1 lines, be aware that the communication from one vault to another is done over a single connection. Unless the links are bonded, only one of the lines will be used for a connection. (Having multiple unbonded links is only of value if there are multiple agents or vaults that will each have separate connections, thus allowing some connections to be made over each of the links.)

*Vault Re-Synchronization.* If one vault loses its disk system and needs to be fully restored from another vault, then in addition to keeping up with the on-going change, there needs to be enough bandwidth to finish the vault-to-vault re-synchronization in some reasonable number of days. One way to reduce the need for extra bandwidth is plan to do the re-synchronization locally by temporarily placing the new vault at the same facility as a functioning vault and doing the re-synchronization over the LAN. The vaults can then be separated, and then the "catch up" time required is only the time to replicate changes that occur while the new vault was in transit to its final location.

If vault re-synchronization will be done over a network link, you should calculate the effective transfer rate between the sites and insure that the time that will be required is acceptable to you in terms of having your backup versions exposed to loss by a second failure during the time that vault recovery is in progress.

**Bandwidth Between Agents and/or Appliances and the Vault.** Remote agents and/or appliances that back up to a vault site need bandwidth for the initial backup and then on an on-going basis for delta backups.

In addition, in doing bandwidth planning you should consider what the restore times will be to a remote site over the site's bandwidth connection.

A rough and conservative rule of thumb is to plan on backing up 100GB of protected data, or less, per T1 line (1.5Mbps). You can scale this rule up or down depending on the actual link speed you have at a site. You may be able to double the "100GB per T1" rule if you have a significant amount of static data, only backup once a night, have a low change rate, or have very compressible data. On the other hand, if you have drawings or image files that don't compress, you may only be able to protect a smaller amount of data per T1.

## **4. Backup and Restore Best Practices**

To plan your initial backups and to set your recovery time expectations, you should have an estimated GB per hour rate for each site. If you already have a connection, the most accurate method is to install the agent software on a server or a desktop at the site. (Windows XP is OK for this.) Do the initial backup of about 5GB without a bandwidth throttle, and see how long it takes.

If it's not possible to test with the actual agent, you can use an estimated GB per hour rate. If you assume close to 100% network efficiency, and you assume that data compresses 2:1, then a dedicated T1 line (1.5Mbps) will allow the movement of about 1GB per hour<sup>7</sup>. To allow for network overheads and the impact of typical bandwidth latency, we recommend that you assume only 80% efficiency. Thus two equivalent ways of thinking about bandwidth are:

- Expect to move .8GB per hour on a 1.5Mbps line or
- Expect to move .5GB per hour for every 1Mbps of bandwidth

If the bandwidth is heavily used with other traffic, the actual throughput will be less.

**Initial Backup.** Iron Mountain recommends that a large server be protected by having several backup policies, each policy protecting no more than 200GB. To do this, create the first policy and protect a subset of the data. Let this policy complete its initial backup both to the TurboRestore appliance (if there is one) and to the vault before adding the second policy and placing more data under protection.

(You can continue to run a tape backup procedure until the initial backup is complete. For instance, the agent is compatible with Veritas and the Veritas open file manager.)

If a policy's expected initial backup time to the vault is more than 60 hours, we recommend that the initial backup be done in stages. To begin, set the backup policy to include a subset of the data with the goal of having the initial backup finish in about 40 to 70 hours. When this initial backup is complete, edit the backup policy to include more data, and so forth until all the data for that server or policy has been backed up.

For large sites, you may elect to "seed" the initial backup through the use of vaults or appliances. Seeding is discussed below.

If the overall expected initial backup time for a server is more than 10 days, the limited bandwidth is likely to be an on-going source of problems. After periods of heavy activity, more than a day may be required to once again have a current backup. You should consider improving the available bandwidth to the vaulting site.

*Getting Started.* For each site, you will want to make a plan as outlined below:

- **Estimate the Initial Backup Time for each server.** Normally, you will not want the initial backup to be running during business hours without a bandwidth throttle. So don't plan on the measured or estimated GB/hr rate being available 24 hours a day during weekdays.
- **Make your plan.** A typical plan is to backup the most important servers at a site over the first weekend, and other servers at night or on following weekends until all the servers and data at the site have been protected. (Don't plan on doing multiple servers at the same time.)

---

<sup>7</sup> Data that does not compress can transfer at ½ GB per hour at best.

- **Use bandwidth throttles appropriately.** A good practice is let an initial backup run over the weekend with no bandwidth throttle. If a server has not completed on Monday morning, create a bandwidth throttle schedule on Monday morning that limits the bandwidth usage during weekday business hours. This way the initial backup will continue during the business day but at a background rate.
- **Don't stop your current backup practice right away.** You should continue with your current backup practice until the initial backup has completed and a few days of delta backups have occurred. There is no reason to place yourself at risk of having no backup for a period of days.

**Seeding the Initial Backup.** If you have large sites, or sites with limited bandwidth that are geographically separated from the vaulting sites, you may want “seed” the backup by using vaults or appliances. This eliminates the need to stage the initial backup over a period of weeks. There are two approaches to seeding:

- Temporarily place a vault at the remote site, and do the initial backup over the site's LAN. Then move the vault to the true vaulting site.
- Place a TurboRestore appliance at the remote site, do the initial backup just to the appliance, then ship the appliance to the vaulting site. The initial backup on the appliance will then replicate to the vault over the vaulting site's LAN. The appliance can then be returned to the remote site, if desired.

**On-Going Backup.** The daily data change rate will have a major impact on the bandwidth needed for on-going backup. If a site falls within the “up to 100GB on a T1 line” guideline, then you don't normally need to worry about bandwidth. If you want to look at bandwidth in more detail, then Iron Mountain uses the following guidelines to estimate change rates when no other information is available:

- 5% for file data
- 10% for database data

Your Iron Mountain Sales Engineer has a tool that will tell you how much of the data has been static in the last month. This may be useful in arriving at an estimate for the daily change rate.

You can apply an estimated change rate to the total data to estimate the amount of data that needs to be sent in a 24 hour period. Then, using the estimated GB per hour rate for a site you can estimate the amount of time that will be required each day for off site vaulting. You normally want this time to be less than the number of work hours (if you are using continuous protection) or the number of hours in the nightly “backup window” (if you are using once a night scheduled protection).

Using a once-a-night schedule usually reduces the amount of data that has to be sent compared to a continuous schedule because blocks of data that change multiple times during the day are only sent once as opposed to many times. We recommend using a once-a-night schedule in cases where the bandwidth is poor.

**Recovery Time.** The time that would be required to fully restore a large server can be measured in days. Even when a full restore is done over a LAN from a turbo restore appliance, the recovery can be long. Let's look at the possible bottlenecks:

- **NTFS.** The bottleneck may be the Microsoft file system (NTFS) on the server receiving the data. An appliance and the LAN are probably capable of delivering files and data more rapidly than the receiving file system can process the data. NTFS performance is not linear as the number of files grows. You cannot restore 1000 files into an empty file system and then extrapolate what the performance would be for a million files. Performance degrades as the file system's internal data structures grow. This degradation is often not noticeable when users create or access individual files, but is very noticeable when the effect is multiplied because millions of files are involved.

Restore time performance will be especially poor if there are a large number of small files as opposed to a small number of large files. NTFS takes much more time to create a file and write the

first block than it does to just write a block into an existing file. The “create file” command is one of the slowest operating system commands.

NTFS performance is especially non-linear with respect to file create times as the number of files in a directory grows. It takes much more time to create the 10,000<sup>th</sup> file in a directory than it does to create the 100<sup>th</sup> file. If a large server has a balanced directory structure where no single directory contains more than 5,000 entries<sup>8</sup> the restore times will be significantly better than a server that contains some very large directories.

- **Hardware.** Obviously, the speed of the hardware on which NTFS is running will have an impact. The speed of the disc I/O system (disks and disk controllers) is particularly important. Memory improves performance because it allows more file system caching. Multiple CPUs will probably not have a significant impact.
- **Bandwidth.** If the restore is done over Internet, then the bandwidth and bandwidth latency time are probably the limiting factor on restore performance. Trying to predict the impact of bandwidth latency is difficult; it's best to run some tests and get actual transfer rates.

*Optimizing Recovery Time.* Some steps that will improve the time to do a full recovery are listed below:

- *Develop a recovery plan.* Some of the data will be more critical to have restored and available before other data. Identify what this data is and thus if a full restore is needed, first issue a restore request that just restores the critical data. Let this restore finish before issuing further restore requests.
- *Do a large restore in stages.* Divide the overall recovery problem into multiple restore requests and let each finish before starting another restore. A good practice to follow is to keep each restore request under 200GB.
- *Be sure there are no excessively large directories.* Plan to keep all directories under 10,000 entries if possible. (Your Iron Mountain Sales Engineer has a tool that will identify large directories.)
- *Use hardware with fast disks and disk controllers.*
- *Use a TurboRestore appliance that meets the TurboRestore hardware spec.*
- *Plan your bandwidth consistent with your recovery time objective.* You should do a test restore of a few large files, several GB at least, to see what transfer rate you actually get. (If you have a TurboRestore appliance, disable it before doing the test<sup>9</sup>.)

## **5. The “Glue” – Ports, Web Servers, Command and Control**

Iron Mountain provides portions of the overall LiveVault Software solution at its own expense and at its own facilities. In particular, Iron Mountain provides the web user interface, the backend SQL database that holds the backup policies, user profiles, logs, and so forth, as well as the command and control system known as the LiveVault Bridge service. This has many benefits:

- You don't have to purchase hardware and software to run this backend infrastructure
- You don't have to install or operate these components
- If you call for Technical Support the Iron Mountain staff can see your backup policies (but obviously not your data) and thus are much better equipped to assist you.
- Iron Mountain provides you with alerts under various conditions. Iron Mountain augments the monitoring you may do of your own IT infrastructure.

The concern with this approach is that if you lose your network access to the Iron Mountain services you would be unable to restore data. To address this concern, Iron Mountain provides a “disconnected restore” interface that you can use locally to restore data from your TurboRestore appliances or vaults when you do not have an outside connection.

---

<sup>8</sup> “Entries” include both files and subdirectories.

<sup>9</sup> Call Technical Support if you need assistance. Be sure to disable the appliance through the interface, don't just turn it off. Turning it off can significantly delay the time the restore job takes because there will be a large time-out period at the beginning of the restore while the agent discovers that the appliance is not available at the moment.

Iron Mountain has two registered ports for the LiveVault software. Agents and TurboRestore appliances communicate to the Bridge service over port 2145. If this port is not open, several other ports will be tried, however, you should plan to open port 2145 for outbound connections so that LiveVault's command and control traffic flows over this designated port.

Agents, appliances and vaults communicate to vaults over port 2144. Other ports will be tried if this port is not available, however, any site where you protect data should allow outbound connections on port 2144.

Vaults will listen for connections on port 2144. Thus any site that has vaults should allow inbound connections on 2144.

## **6. Vault Monitoring**

Vaults need a significant amount of temporary disk space during the times that backup versions are being merged together. A vault should always have 30% or more free space. You should monitor this and when the vault approaches 60% full you should provide more free space by:

- Adding more disk space to the vault. All the disk space available for vaulting needs to be in one large partition.
- Off loading some of the agents that back up to the vault to a different vault.

You should limit the number of agents that back up to a vault to 100 or less if they are using continuous protection.

## **7. Conclusion**

Please discuss your particular environment, backup and archiving needs with your Iron Mountain Sales Representative and Sales Engineer. We are here to help you create the best solution for your organization. Iron Mountain also has a Professional Services team that can work with you in greater depth to design and implement your solution.