# Important Considerations When Selecting a Branch Office Data Protection Solution

Organizations such as retail chains, schools, hospitals, insurance companies, and many other businesses with remote locations face unique challenges when it comes to protecting their mission-critical data. Ensuring that business-critical data located in geographically dispersed sites throughout the organization is protected regularly, and in accordance with established backup and recovery policies, requires a comprehensive approach.

This paper discusses the top considerations storage administrators need to keep in mind when planning, selecting, and installing such a Branch Office Data Protection (BODP) solution.

Assume that your company has multiple remote sites, each with a heterogeneous network of computers running a variety of operating systems including Windows, Mac OS X, Linux, VMware, and UNIX. The sites are connected by VPNs over the Internet. The sites have little or no IT personnel at the remote sites, so the Branch Office Data Protection solution must be capable of running continuously and reliably, while administration and data protection takes place at the central office.

## 1. Right-sized Storage for Each Site

To ascertain how much storage you will need to back up your local NAS system at the site, determine the amount of storage used by those computers and multiply that amount by 1.5 to 2.0 to conclude how much storage is needed to contain your backup data. The total storage required is equal to your primary storage needs plus the additional storage required for backups. Storage needs often grow by 80% per year, so be sure to plan for your future needs, as well as what is required today.

Once you've determined your storage needs for each site, purchase the storage system that best fits the needs of each location, being careful not to over purchase. Since the initial purchase price of a BODP solution is heavily dependant upon the size of the systems you buy, finding a vendor that offers systems in a variety of sizes – from small desktop units for your smaller locations, to scalable rackmount systems for your central office – is key to obtaining the storage infrastructure you need, without breaking your budget.

## 2. Backup Servers

Larger sites typically have a backup server to collect the backup data from each computer and copy it to disk or directly to tape. However, the cost and ongoing maintenance of a backup server can be prohibitive for smaller sites. A significant alternative for these sites is to use the NAS system, itself, as the backup server. If the NAS system is powerful enough, it can run the backup software directly on itself to increase efficiency and save cost.

## 3. Protect Everything

Desktop and notebook computers often contain up to 80% of the company's business-critical data. Therefore, it's essential to protect them by backing them up to a NAS system whenever they are connected to the local network. Desktop and notebook users should be able to perform ordinary restores themselves without the aid of IT personnel. Heterogeneous file servers and application servers can be backed up either directly to the NAS system, or via a backup server which places the backup data onto a NAS system using the CIFS, NFS, or iSCSI protocol over Ethernet.

## 4. Link Speeds Matter

All too frequently, the speed of a site's incoming telecommunication link is frustratingly slow. A T1 line is most commonly used, but that only provides a maximum throughput of 1.544 megabits per second, which translates to a maximum speed of 1 GB per hour. Conversely, a T3 line provides 28x the maximum throughput of a T1 – but at $120,000 per year, this option is cost prohibitive for many sites.

## 5. Onsite/Offsite Data Protection

The typical pattern for backups is to do a monthly full, weekly full, and daily incremental backup to tape. After a backup is completed, it must not be stored exclusively onsite. Your data should never be considered fully protected until a copy is stored in an offsite location. A best practice is to duplicate the backup tapes, then keep the originals onsite for rapid retrieval, while storing the duplicates in a safe offsite location.

It is well known that backing up to tapes has several problems. Tapes degrade with age and wear out with repeated use. The tape drives, heads and mechanism wear out as well, and can be costly to replace. Tapes must be managed, secured, stored, and replaced. A complex schedule often referred to as "grandfather-father-son" tape rotation needs to be implemented. The task can be overwhelming for a remote site to properly perform.

Alternatively, backups can be stored on disk at each site for fast backups and restores. They can then be transmitted from the remote site to the central office over telecommunication lines. However, a 25 GB backup will take over 25 hours to transmit over a T1 line. While daily incremental backups may be small, the weekly and monthly full backups can generate an extraordinary amount of backup data, which can inundate telecommunication lines.

To circumvent this problem, each site should perform their first full backup to onsite disks or tapes, then copy that initial full backup to the central office for safe keeping. All subsequent backups are incrementals, consisting only of the bytes that have changed. In this manner, they remain small enough to be transmitted electronically each day to the central office.

Additionally, a virtual full backup can be fabricated for any day from the full and incremental backups that have been taken. Consolidations eliminate restore points to save storage.

## 6. Replication and Backup and Restore Software are Complementary

The secure replication of data between sites is an essential component to your BODP strategy. Replicating business-critical data from each of your remote sites to the central office ensures that nothing is left to chance. The essential data from each site can be treated consistently by qualified IT personnel in accordance with your established backup and recovery policies.

Following the initial replication between a server at a remote site and a server at the central office, only the changes need be transmitted. Such a replication is often run continuously. The first replication is often done while the two computers are co-located to speed the process.

The weakness of replication is that if a virus, human error, or other disaster affects one of the servers, the problem is faithfully replicated to the other. To guard against this weakness, the creation of additional restore points is highly recommended. Trained IT administrators can perform daily backups of the replication target in the central office. This practice creates a stopgap of the data loss, since the previous day's backup can be restored in the worst case scenario.

## 7. Using a Virtual Tape Library

To overcome the problems associated with backing up directly to tapes, it is often desirable to back up to disk as a more efficient alternative. However, many backup software applications only know how to write to tapes to perform full and incremental backups. Virtual Tape Library (VTL) software bridges this gap by making a disk look like a standard tape library to the backup software. VTLs don't stretch and you don't have to remember to load the tape each day. The best practice is to make physical copies of the virtual tapes and store them offsite. Some backup software allows the virtual tapes to be transmitted to the central office, where they can be committed to physical tapes by experienced backup administrators.