

## Remote Disaster Recovery

### **INTRODUCTION**

Remote disaster recovery sites are now feasible and affordable for small and medium businesses. Once the province only of major Wall Street brokerages due to the phenomenal expense of mainframes, software and private leased lines, the cost of today's technology enables virtually every business to implement a cost-effective remote disaster recovery site using off the shelf hardware, software, security and communications. Even if you do not have an alternate site, many off site vaulting locations are available for a reasonable monthly fee.

#### *Replication frequency*

The replication strategy depends on the data to be backed up. The more valuable the data and the more irreplaceable, the more frequent the replication. There are two types of replication synchronous and asynchronous. Synchronous replication over IP networks would entail unacceptable delays in the execution of the primary application and is limited to networks with private leased lines and large financial transactions. Asynchronous replication queues up the data and transmits it, as bandwidth is available. It is by far the most practical for the vast majority of applications. There are two popular types of backups that are replicated via asynchronous replication below.

#### *Snapshots*

Taking a "snapshot" backup of your data is the easiest way to perform a remote replication. After a snapshot is taken, the incremental changes are sent to the remote site. There is a tradeoff between frequency and bandwidth. Data rewritten multiple times between snapshots is transmitted only once with the final data. Typically snapshots are taken several times per day. Databases need to be quiesced momentarily during the snapshot and special database agents are provided for applications like Oracle and Microsoft Exchange. This solution is best for recovering from a local disaster with loss of only a fraction of a day's work.

#### *Continuous Data Protection*

A more advanced solution for replication to remote sites is known as continuous data protection. In this scenario, all of the data is time stamped so that the system can be rolled

back to any point in time called the recovery point prior to a problem or catastrophe. This is especially useful for recovering from a virus, or rolling catastrophe where an event occurred and slowly but surely corrupted data. In the event of a local disaster, the data can be rolled back virtually to the moment of the event with little or no loss of data.

#### *Operational Recovery*

The disaster recovery site can be optionally configured with servers to stand in immediately for the primary servers so that the recovery site data can be placed into service promptly. This protects the business operation as well as the data.

### **WHY REMOTE DISASTER RECOVERY SITE**

- Protect primary site
- Reliable
- Simple to set up
- Low Cost

### **REQUIREMENTS**

- Alternate site
- IP bandwidth
- Replication software
- Disk arrays
- Replication appliance/server
- Optional replacement servers

### **TYPICAL ENVIRONMENTS**

- High cost of data replacement
- High cost of downtime

### **BENEFITS**

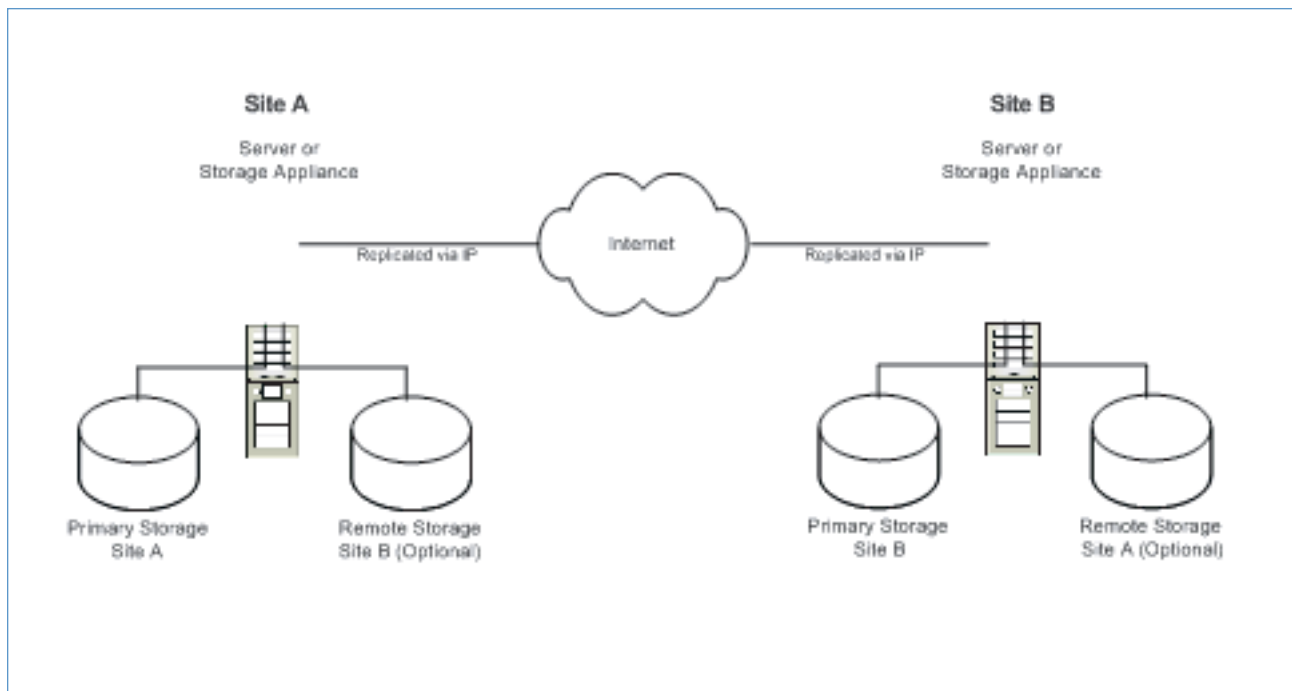
- Data protection
- Disaster recovery

## HOW IT WORKS

Data recorded at the primary site must be initially duplicated to the secondary site. Mirroring the data to disk locally and physically moving the duplicate copy to the remote site is the most typical way to perform this initialization. Alternately, data is copied to tape and then copied to the disks at the disaster recovery site.

Once initialized, only changes need to be transmitted to the disks at the disaster recovery site. The system administrator determines the frequency of transmission of each type of data to be copied to the disaster recovery site with the software tools being used. To obtain the best results for the least cost, skilled storage architects carefully match your specific requirements to the software capabilities and options to configure a disaster recovery solution.

Once specified, the host based software or network based appliances automatically follow the system administrator's policy directions. The system administrator can change performance, storage and cost considerations at any time.



**In this illustration, two sites are mutual disaster recovery sites for one another. Each site mirrors data to the other site. In the event of a disaster at one site, the other site has up-to-date information. Using standard IP technology, each site may be located anywhere in the world. Optional replacement servers at the remote site can stand in quickly for the primary site.**