



## Remote IP-based Replication Option – At a Glance

### What is it?

The IPStor Remote IP-based Replication Option provides an IP-based mechanism for data replication from one site to another, for purposes such as off-site disaster recovery, site migration, electronic vaulting, etc.

### How is it used?

Specifically designed to defend against site failure by providing automated off-site data protection, the Remote IP-based Replication Option provides fast remote data synchronization (FRDS) of data on disks managed by IPStor software (virtual or service-enabled<sup>1</sup>) between one IPStor appliance and another—across the street, across town, or across the globe. In case of a catastrophic failure at the primary site, the system administrator can quickly redirect application servers to access data from replicas located in the backup data center.

Administrators can specify a variety of policies to control the replication process, giving them a very granular policy-driven mechanism for keeping an extra set of data off-site for disaster protection. The Remote IP-based Replication option makes use of IPStor software's advanced snapshot technology, which protects data from long-distance transmission problems and guarantees the integrity and usability of replicated data. Therefore, if a problem occurs during a subsequent replication, the IPStor software will use the data from the Snapshot Resource to recreate the replica from its last *good* state.

As part of FalconStor's Continuous Nearline Backup solution, the Remote IP-based Replication option works in concert with the TimeMark Option<sup>2</sup> to deliver disaster-protected full, incremental, or differential policy-driven backups to low-cost disk, without backup software. The Remote IP-based Replication option provides this extra layer of protection by enabling delta snapshots to be taken of replica data at a remote site (centralized data center or disaster recovery site). There is no impact to application servers, making "on-peak" backups of live data possible; the issue of the backup window is altogether eliminated and recovery of individual files/records or volumes is immediate.

### How does it work?

After a disk managed by IPStor software is created, the storage administrator – using the IPStor Console – designates a replication target volume at a secondary, remote location. The IPStor Appliance at the primary site (source) communicates with the IPStor Appliance at the secondary site (target) to create the replication channel. An initial synchronization process must be performed prior to commencement of scheduled replications. One method will synchronize the primary and replica disks as soon as the initial replication configuration is complete. An alternate method, available for large data sets, scans both the primary and replica disks to determine the amount of data (block-level differences) to be synchronized. Once the scan is complete, the synchronization will occur at the next scheduled replication. After the initial synchronization is complete, only changes to the primary data are sent over the wire.

Replication can be scheduled on a per volume basis according to various policies, such as time of day, time increment, or amount of data that has changed. The IPStor software acknowledges the write requests to the

---

<sup>1</sup> Enables existing disks to make use of the full suite of IPStor storage services, including Mirroring, Replication, Snapshot Copy, TimeMark, HotZone, \_SafeCache, etc. Original disk partitions/volumes are not changed; no data conversion or migration is necessary.

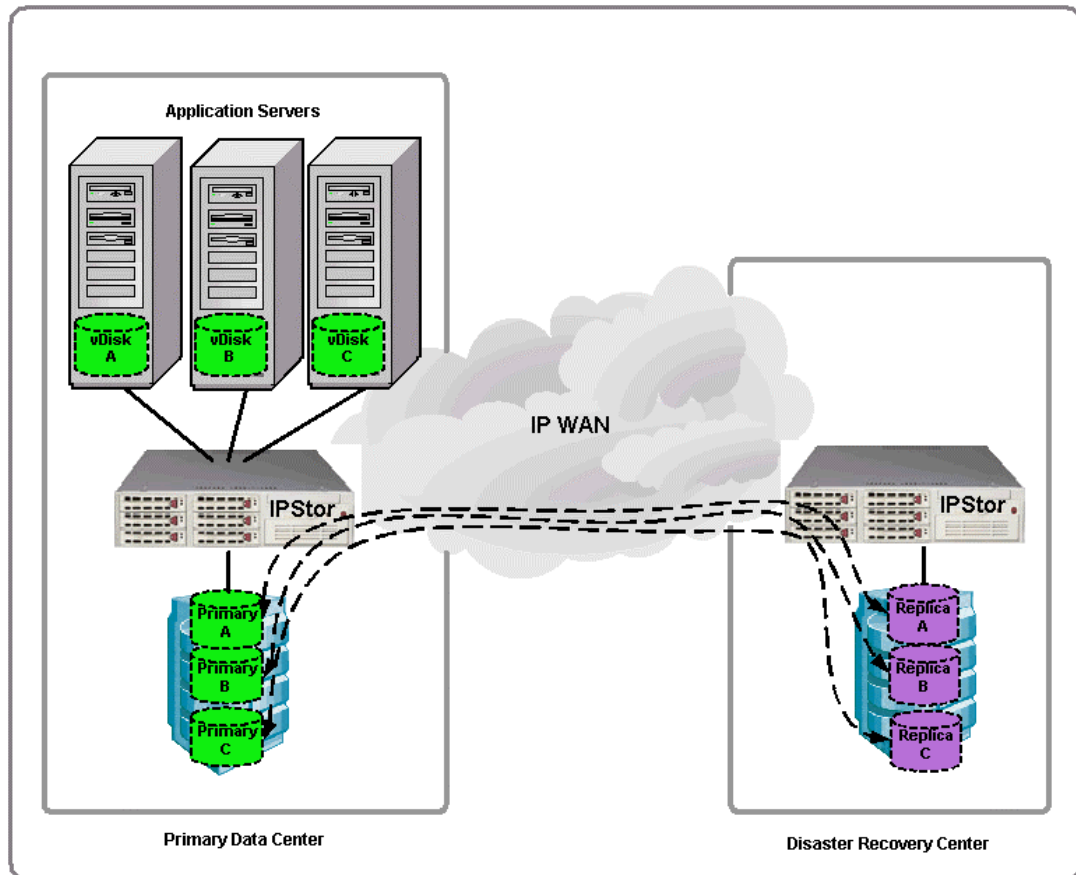
<sup>2</sup> For immediate recovery of files, records, or entire volumes lost to accidental deletion, corruption, viruses, hackers, etc., the IPStor TimeMark Option enables administrators to create multiple (up to 256), instantaneous point-in-time snapshot copies - "delta snapshots" - of an active drive. These delta snapshots can be created automatically based on an administrator-defined schedule, for ease of use and high reliability. Administrators can then instantly recover individual files or volumes by mounting a delta snapshot as an active volume by using the included TimeView™ feature.



application server as soon as they are written to the primary disk. The off-site data movement occurs independently and has no impact on application performance.

For example, when a watermark is reached, a snapshot isolates the data to be replicated. The snapshot creates a point-in-time image that is used as the source for the replication, ensuring that all data copied is valid for a given point in time, rather than copying data as changes are being made. Use of the IPStor Snapshot Agents for Databases or Messaging Systems ensures transactional integrity for database/messaging applications.

Should there be a storage failure at the primary site, the replica volumes are promoted into fully usable volumes. These can then be accessed via IP, allowing immediate access to the data. Resumption of business can now be measured in minutes, rather than days. When the primary storage is brought back online, the IPStor software can re-sync the volumes – again by sending only changes over the wire, not all of the data – allowing the primary site to receive all data accrued while it was down. The storage network returns to its original state, and the replication cycle returns to its normal patterns.



This diagram illustrates the primary IPStor appliance communicating to an off-site IPStor appliance via a standard WAN connection for replication from one site to another.

### How does this benefit me?

#### Reduces hardware/software costs for Replication

With the IP Replication Option, data is replicated from one IPStor Appliance to another IPStor Appliance over any existing LAN, MAN, or WAN network infrastructure without the need for extra FC-to-IP converter boxes. The



© 2003 FalconStor Software. All Rights Reserved.  
125 Baylis Road  
Melville, NY 11747  
631.777.5188

[www.falconstor.com](http://www.falconstor.com)

IPStor Remote IP-based Replication option is performed at the IPStor appliance and is independent of application servers and operating system platforms. There is no need to deploy matching file servers or application servers at the disaster recovery site (except for hot standby scenarios). As a software-based, vendor-neutral solution, IPStor offers unprecedented flexibility in creating disaster recovery environments. Because IPStor software replicates between IPStor Appliances, the source and target storage hardware need not be the same, allowing for low-cost Disaster Recovery (DR) planning by using low-cost JBOD at the DR center. Multiple disks managed by the IPStor software in various locations can also be replicated to a single target storage device.

**Delivers disaster-protected instantaneous incremental backup of live data to low-cost disk with easy, immediate recovery**

For robust data protection, the IPStor Remote IP-based Replication Option works in conjunction with the TimeMark Option to provide disaster-protected full, incremental, or differential automated instant backup to disk. Now administrators do not have to spend needless hours trying to recover accidentally deleted records or virus-infected files from tape, even in the event of a disaster.

**Simplifies and accelerates the entire process of pre-and-post disaster operations**

Before the disaster, the initial synchronization can be done using tape backups or mirroring. A delta-sync process facilitates minimal transfer of data over the WAN. During an emergency, IPStor's SAN-over-IP protocol driver can be deployed to allow emergency access of the data at the DR center over WAN by any server located anywhere. This IP connectivity provides significant advantages, offering many more ways for temporary offices to access data during an emergency. A reverse-delta-sync process facilitates fast recovery of the primary site when the emergency is over.

**Snapshot Agents minimize recovery time**

When used in conjunction with the IPStor Snapshot Agents, the replicated data has full transactional integrity in addition to point-in-time consistency. This means the replica can be immediately put into active use, without going through a complete 'consistency check' process that can be very time consuming for large databases.



© 2003 FalconStor Software. All Rights Reserved.  
125 Baylis Road  
Melville, NY 11747  
631.777.5188  
[www.falconstor.com](http://www.falconstor.com)