

iSCSI Storage Server

White Paper


iSCSI Storage Server is a SCSI/Fibre Channel network disk controller intended for use with Bus-Tech's Mainframe Appliance for Storage.

This paper introduces the reader to the concepts of Mainframe Appliance for Storage and the iSCSI Storage Server and then reviews the alternate ways the iSCSI Storage Server can be implemented within the Mainframe Appliance for Storage environment.


Overview

As shown below the Mainframe Appliance for Storage (MAS) is a tape-on-disk controller for IBM System/390 class mainframes. MAS connects to the mainframe using 1 or 2 ESCON channels and emulates 1 to 32 3480 tape drives. Tape volumes written to any of the MAS tape drives are actually written to disk with each tape volume becoming an individual disk file.

By itself the MAS supports connection of several different types of disk storage including SCSI (ultra160), fibre channel, or network attached storage. Network storage is attached via a gigabit Ethernet interface on the MAS. The MAS supports both Network File System (NFS) and Common Internet File System (CIFS) protocols for communicating with Unix, Linux, or Windows servers as well as commercially available NAS devices. The MAS also supports the iSCSI protocol for connection to the iSCSI Storage Server.



iSCSI Storage Server



Bus-Tech's iSCSI Storage Server (ISS) is a SCSI (ultra160) / fibre channel disk controller designed to export disk volumes over an IP network using the iSCSI protocol. ISS was specifically developed for deployment with Mainframe Appliance for Storage to offer a high-performance option for remotely attaching storage over an IP network. iSCSI device drivers are included on every MAS shipped allowing it to communicate with the ISS. Logical units (LUNs) exported by the ISS appear to the MAS as any other SCSI or fibre channel attached disk would appear as opposed to NFS or CIFS attached storage which requires the MAS to use appropriate network client software.

There are four primary uses of the iSCSI Storage Server with the Mainframe Appliance for Storage. They are:

- ✓ As the MAS primary storage
- ✓ As a backup device to MAS primary storage
- ✓ As a synchronous mirror to MAS primary storage
- ✓ As an asynchronous mirror to MAS primary storage

The remainder of this paper will take a look at each of these uses in detail.

Primary MAS Storage


When the ISS is used to provide primary storage to the Mainframe Appliance for Storage, the MAS will normally have no local (SCSI or Fibre Channel) attached disks. Instead the MAS connects to its storage via the iSCSI Storage Server using gigabit Ethernet, IP, and iSCSI.

There are several reasons why connecting storage in this fashion may make sense. For example:

1. No SAN is in place and you need to place storage at distances greater than SCSI channels will support.
2. An existing SAN is in place but not within distance of the MAS. A corporate IP network is available to bridge the distance.
3. You would like to share a SCSI RAID storage unit among more than one MAS. By dividing the RAID storage in to several logical units (disk volumes / partitions) the iSCSI Storage Server can export a different volume to each MAS controller.

iSCSI Storage Server

Backup Device to MAS Primary Storage



The iSCSI Storage Server makes an excellent alternative for backing up the MAS primary tape library. The MAS provides standard dump / restore routines allowing the MAS to periodically dump incremental or full backups of the tape library to backup datasets.

Use of the ISS to provide storage space for housing dump datasets improves performance of daily dump processes and keeps backup files online and ready to use at a moments notice. Restoring a lost or damaged tape volume becomes a matter of interactively running the MAS restore routine at the MAS console and selecting the tape volumes that need to be restored. There is no need to locate, retrieve, and/or mount backup datasets because they are already online and available.


Further, since your backup storage is IP-attached, routers and/or network gateways can be used to place your iSCSI storage server at a remote location or disaster recovery site.

Synchronous Mirror

One of the real advantages of iSCSI-attached storage when compared with NAS-attached (NFS/CIFS) storage is that iSCSI storage appears to the MAS Linux operating environment as a real SCSI device rather than a networked device. This difference means that the file system used to read/write from iSCSI storage is local to the MAS instead of the NFS/CIFS model where the real file system performing input/output to disk is located on the network attached storage server.

One of the benefits looking like a real SCSI disk implies is the ability to perform synchronous mirroring. The MAS, which is a Linux-based controller, includes the ability to provide a software-based Redundant Array of Independent Disks (RAID) system. RAID-1 (mirroring) maintains duplicate copies of data on two equally sized SCSI disks.

iSCSI Storage Server



As shown above, software RAID-1 allows the MAS to continuously maintain duplicate, mirrored copies of the tape library at all times. When the mainframe issues a write to one of the MAS tape drives the MAS actually issues two writes, one to each mirror image. MAS will not acknowledge completion of the write to the mainframe until both the local (SCSI/fibre channel) and remote (iSCSI/IP) disks have completed writing the data.

If either of the disks should fail, the MAS RAID software registers the unit as having failed, but keeps operating using the remaining disk. The mainframe is not aware there was a problem but the MAS operator console is informed that the RAID is running in reduced operations mode.

When the failed disk unit is returned to service the MAS will rebuild the out-of-date mirror and full operations will resume. MAS emulated tape drives never interrupt service to the mainframe.

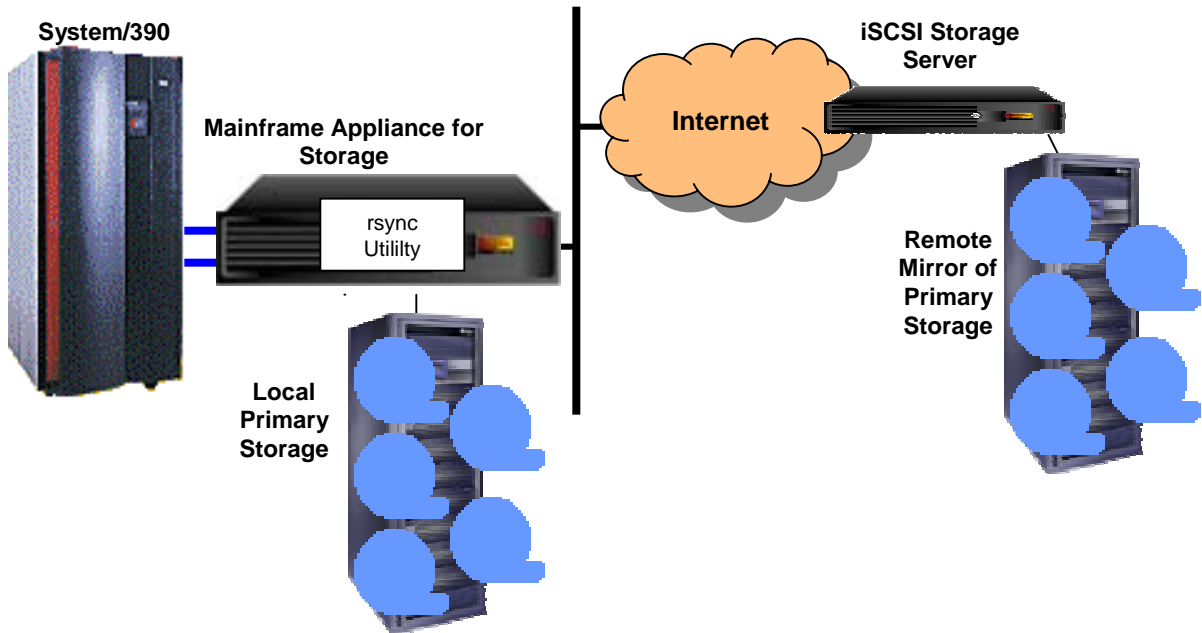
Asynchronous Mirror

Perhaps instead of maintaining a RAID-1 mirror using high-performance networking, you would like to maintain a remote mirror of the MAS library periodically updated using the Internet as a transport mechanism.

Software RAID is inappropriate because you cannot be guaranteed of adequate throughput or response to prevent time outs of the remote writes.

The MAS provides the Unix/Linux rsync utility as part of its collection of backup/networking utilities. Rsync will compare two directories and update the destination directory so that it exactly mirrors the source directory.

iSCSI Storage Server



When using the iSCSI Storage Server to connect a remote disk to a Mainframe Appliance for Storage, rsync can be periodically run in order to make sure the remote copy of the tape library is a point-in-time copy of the local tape library.

If the local tape library becomes disabled the remote iSCSI-attached copy can be used to begin processing in reduced operational capacity. Only those tape volumes updated since the last rsync execution will need to be recreated. And in the unfortunate event that the entire local copy of the tape library is lost, rsync can be used in reverse operation to retrieve the library across the internet from the remote backup site.

Summary

The Bus-Tech iSCSI Storage Server is a network disk controller capable of exporting SCSI and/or fibre channel disk volumes out over an IP network using industry standard iSCSI protocol.

The iSCSI Storage Server has been specifically designed and implemented for use with Bus-Tech's Mainframe Appliance for Storage; providing MAS with iSCSI network-attached storage space.

Customers implementing Mainframe Appliance for Storage can use the iSCSI Storage Server to provide a range of enhanced capabilities to the MAS solution including simple IP-based channel extension, local and remote backup, synchronous mirroring of the MAS tape library, and/or remote data vaulting for disaster recovery.