



Backup Appliance – A Definition

Every business operation today has data management requirements around the three data protection areas of backup, archive and disaster recovery. A Backup Appliance should address all three areas:

- **Backup** – on-line copy of every file a customer has identified for restore. The default should allow for the standard business requirements of a 45 day retention of any file once it is deleted or changed. A customer, however, can set the data retention time for a file at any length, shorter or longer. Backup files should include all files – op/sys, hidden files, databases (DB), flat files, DOC/XLS/PPT files, etc.
- **Archive** – on-site point-in-time capture of any file or set of files. These files need to be saved for some business or legal time requirement, from months to years. Also, these may be files that no longer need to sit on more expensive active or spinning disks, but need to be moved to a cheaper media.
- **Disaster Recovery** – copies of all on-line copied data (backup) and on-site saved data (archive). Customers can change their disaster recovery copies to one or more locations, preferably copied simultaneously. Off-site copies can be media removed manually, or copies electronically vaulted. A daily Disaster Recovery Plan can be optionally provided as a function of DR.

In addition, a backup appliance must allow for automation and long-term management flexibility around:

- **Media** – since media specifications are constantly changing for disk, tape and optical technologies, a backup appliance should provide for virtualization of data storage, allowing the media in a storage pool or bucket to be changed, migrated, updated, distributed and consolidated across locally attached, network attached or SAN managed locations. Over time, the data on media should not be limited by media decisions, and should not require later restoration in order to be moved to a new media.
- **Platforms** – IT departments have to remain flexible in managing their data across many platforms and many networks. The backup appliance should allow for backups and archives to be performed in a like manner across all popular platforms.

Finally, a backup appliance must include all the hardware, software and integration of a backup solution in one appliance. Optimally, the appliance should include a logical plug and play mix of the hardware components for both:

- **Scalability** – a backup appliance must offer a customer the flexibility to grow, both by upgrading the solution they own, and a forklift upgrade option to move seamlessly into a larger solution.
- **Support** – with all the hardware, software and integration for a backup solution included in one delivered product, the backup appliance should provide for a one warranty, one support maintenance capability.

By definition, a backup appliance must plug into a data environment through whatever network path is most logical for data movement, and provide a true “enterprise” capability. Users should be backed up, whether they are on a LAN, WAN or whether their data is located locally or in some virtual NAS or SAN location:

- **LAN friendly** – in circumstances where data must be copied off a LAN location over a fiber or other backbone, the backup appliance must operate without problem, but most data should be handled over the common traffic area used by the bulk of users, without additional burden to the network. As networks increase in speed, the backup appliance must operate without burdening the normal traffic patterns on the LAN. In the case of WAN users, options for byte-level change backups must be provided.
- **Data agnostic** – data can be located by the backup appliance on any of the many locations where users access files, including locally attached, mapped drives, NAS locations, SAN locations, and remotely managed data over intra or internets.

User:

Backup, archive and disaster recovery begin with ease of use from the customer, who will include any person who manages data. This includes individual users, system managers, database administrators, etc. Also, a backup appliance must apply services to all systems, including laptops, workstations, servers, clusters, etc.

- **Easy to use** – minutes a day to manage and administrate – simple to navigate
- **Plug and Play** – installs in minutes – network connectivity, component connectivity, and client connectivity
- **Heterogeneous** – backs up and restores all popular platforms. Clients can be added to the system easily and remotely

Functionality:

Backup will move data to an on-line location where all files can be restored that a customer identifies as meeting a restore requirement. This on-line location backup should be managed by policies based upon a data retention set by a customer, not by limitations of hardware. If customer wants expired or deleted files available for up to one year, they can have that on-line.

Restore functions allow users at any level (from laptop to large RISC systems) to bring back files by communicating directly with the backup service or server, if the requirement is called for. Restores should be available at the file level or some backup set level, both identified by a user.

Archive is a point-in-time capture of any set of active data a customer requires to save for any period of time. Can be kept on-site. Retrieve is the ability to bring back any set of data described as an archive.

Users should have access to their restores and retrieves, if allowed, without the requirement to contact a system manager. This should be an easy to use function of their workstation utilities and tools.

Technology:

Backup should be quick enough to fit any window requirement that can be met by hardware specs

- Scheduled or journal or polling capabilities optional
- Backup media is virtual or “pools” allowing disk, tape, or optical
- LAN, WAN support
- SAN, NAS support
- Backup over any IP pipe

Restore should be quick enough to fit any window requirement that can be met by hardware specs

- File, directory, wild card, disk, or system restores available
- Restore to original location, or another location specified
- Restore by file, wild card, etc.
- Restore by user, manager, system manager, or various password options
- Restore over any IP pipe

Archive should allow data to be kept over any length of time

- Scheduled capabilities optional
- Archive media is virtual or “pools” allowing disk, tape, or optical
- LAN, WAN support
- SAN, NAS support
- Archive over any IP pipe

Retrieve (restoring an archive) should be quick enough to fit any window requirement that can be met by hardware specs

- File, directory, wild card, disk, or packet of files retrieve available
- Restore to original location, or another location specified
- Restore by file, wild card, etc.
- Restore by user, manager, system manager, or various password options
- Restore over any IP pipe