RevStor
To Protect and Store

sales@revstor.com
800-763-4170

The need for alternative storage methods

Data has grown wildly beyond our control. Issues such as data duplication, compliance and regulatory issues have catapulted data to an even more critical status for all companies. The need for easy-to-use, manageable and scalable solutions has reached critical mass, yet most solutions are steeped in the traditions and hence problems of the past.

In parallel with the ever-increasing need for data management has been the ever-decreasing cycle of hardware improvement at even the common desktop level. Today's desktops are orders of magnitude more powerful than even last decade's servers, yet sit idle and unused to even a fraction of their potential.

Enter the technology of mesh networking. Mesh networks harness the power of the whole network to harness the power of the desktops, laptops and even servers in an organization to create a secure, powerful and extremely reliable storage system.

Who can implement mesh storage?

The LAN is one of the most robust and resilient pieces of a companies IT infrastructure. The mean downtime for switches and routers is a fraction of that for servers, desktops and even laptops. Mesh storage rides on this area of the infrastructure, allowing full use of a company's internal resources and balancing the network load more appropriately for storage-related requests.

An ideal implementation has a minimum of between 100 and 200 computers on the network at least 70% of the time and a minimum of 2-3 locations. Scalability of mesh storage (as shown by the consumer counterpart products like Napster and Gnutella) is to the millions of nodes. As the number of computers (nodes) increases so does the resilience and security of the storage system.

Why should we implement mesh storage?

1. As a company with an existing general purpose computer and network infrastructure, mesh storage requires no new hardware purchases.
2. Mesh storage technology also provides full use of the resources the company has already purchased and extends the life of those resources well beyond normal limits.

sales@revstor.com
800-763-4170

3. Mesh storage provides direct control over not just a company's data, but its personnel costs as the nodes are self-managing and the system as a whole is self-healing requiring less maintenance and monitoring time be spent to unsure uninterrupted up-time.
4. By taking part of the IT infrastructure out of the "raised floor" environment multiple costs are eliminated or contained from real estate to capital budget requirements.

How does it work – why should I trust it?

Mesh storage works by taking the unused portions of a computer's resources and pooling those resources with the other computers on the network. A computer's resources consist of hard drive space, network bandwidth and CPU cycles. By combining these resources together, mesh storage works with the network, not against it, to provide its data repository.

Security is of the utmost concern in mesh storage as the desktops and laptops are notoriously unreliable and unsecured. To accomplish this needed level of security, mesh storage must enable certain features:

1. Files must not be stored as a whole when at rest.
2. File pieces must be obscured, encrypted or both in order to prevent even a piece of a file from being viewed by those not approved to do so.
3. The (t, n) threshold for retrieval of data must be a nearly unlimited number outside the network, but effectively 0 inside the network.
    a. In order to retrieve a file a computer must be a part of the network on which the mesh storage system is implemented.
4. In order to ensure reliability of the data, multiple copies of each piece must be maintained such that the system can self-heal in the event of node loss.
5. Management of the system must be able to occur from any node to prevent the loss of management ability due to a particular system failure.
6. Movement of pieces over the WAN must also be available in order to ensure that should a single location be unavailable, that location's data is still maintained.
7. File pieces must be stored in such a way that the potential for piece corruption is eliminated by creating values identifying the piece which are stored independently of the piece.

The key to mesh storage functionality is the intelligent node. Each node in a mesh storage system must be able to act both independently and as part of the group. Nodes are naturally altruistic in design, providing resources to the storage system. However, nodes must also be cognizant of the other users of their resources, namely the User. Nodes monitor resource utilization to reduce and even eliminate impact to the User. Nodes are also capable of deciding, as a group, where to best place data pieces and where to best retrieve those pieces from. This intelligence is a central pillar to any mesh storage system.

sales@revstor.com
800-763-4170

This combination of reuse, security and intelligence provides companies with the most cost-effective storage method available. The ROI on a mesh storage system is months not years and the maintenance of such as system is already defined by current computer and network budgets.

SANware – an example of a successful P2P storage architecture

SANware implements the above features in order to provide a completely heterogeneous mesh storage system for the backup of unstructured data. SANware was designed with both the benefits and limitations of mesh architectures in mind to make best use of the available resources for companies.

Backup was chosen due to the tolerance for latency and the lack of high-transaction read-writes. Mesh storage systems are naturally latent as nodes must be sought out and queried for their effectiveness for participating in a given storage event. This planned latency is necessary for the use of a general purpose infrastructure as outlined above. Though there are methods for reducing latency in node groups, many of these add overhead to both the node and the system as a whole.

SANware is implemented through agent-based technology. Software agents are deployed on the computers designated to be part of the mesh architecture. Those agents the communicate using lightweight protocols such as UDP messaging to ensure low network impact and yet provide near-constant communication. Timing systems are built in to prevent the flooding of the network especially as the number of nodes grows.

In order to successfully manage large numbers of nodes, a natural hierarchy is built in to the storage system. Supernodes are dynamically promoted from within the node infrastructure in order to provide traffic control and minimize node broadcasts. Each node is paired with a supernode and that supernode is also made aware of the other supernodes within a given installation. Supernodes are promoted based on resource availability in order to ensure low impact on Users and the network.

Nodes communicate with their Governing Supernode in order to determine where to send and retrieve data pieces. Nodes also send their resource statistics to their supernode for aggregation and dissemination to other supernodes. These aggregated statistics provide enable multiple functions:

1. Determination of future supernodes
2. Determination of the best storage locations
3. Determination of potential problems within the node groups

As files are sent to the storage infrastructure they are first pieced then hashed based on the piece contents. Those hashed pieces are then compared to other hash values currently stored to determine if the piece already exists. This enables a basic form of deduplication to occur and reduce storage use and growth. The pieces are then encrypted using a

sales@revstor.com
800-763-4170

randomly generated key, duplicated the user-defined number of times to ensure availability and then sent to the most appropriate nodes for storage.

The hashes, keys and other values describing the piece – the piece metadata – are stored in the Distributed File System or DFS. The DFS resides on all supernodes and hence is also distributed. Changes are made to at least 2 copies of the DFS simultaneously, and then replicated to other supernodes in order to ensure that replication does not break down due to the loss of a supernode during the replication process.

Replication of both files and the DFS follows a tree format where node 1 replicates to nodes 2, 3 and 4 and those nodes replicate to 3 additional nodes each until the desired availability value is achieved.

The future of storage?

As computers and networks grow in power, the need for dedicated systems to perform basic tasks decreases. Mesh storage is an outgrowth of that decreasing need. Many of the issues proscribing what a mesh storage system can achieve will change as computers become more powerful and networks less latent. Mesh storage provides a way for companies to bring cost-effective, resilient storage outside the glass walls of IT.