TARMIN

GridBank™ Active Archive

Archive, manage and discover critical business information

The challenge for organizations to efficiently store and manage data is intensifying with the enormous growth of unstructured data.
But the sheer volume of data alone is not the full extent of the problem; it is also the number of locations where this information is stored — on Tier 1 disks, departmental file servers, backup tapes in offsite storage facilities, employee laptops and so on. This makes it almost impossible for organizations to know what data is located where, making retrieval of data and application of compliance and governance policies extremely difficult.

To address these challenges, Tarmin has developed GridBank, an integrated software solution consisting of an intelligently managed searchable archive and next generation secondary storage platform. The storage platform virtualizes all underlying secondary storage infrastructure creating a single logical view across multiple sites and heterogeneous storage devices. This single interface for applications, users and administrators into the secondary storage tier reduces management complexity and allows easy application of ILM policies. An end user search engine allows easy search and retrieval of data, improving functionality for end users and simplifying e-discovery. GridBank delivers several key advantages:

- Locating data is as easy as executing a search on the internet
- Data is always stored securely and compliantly
- Data is automatically managed by policy throughout its lifecycle

Reduce Storage Costs
GridBank is designed to deliver the most efficient storage of unstructured data, reduce overall storage footprint and deliver significant CAPEX and OPEX savings. Open standard interfaces allow data migration from high cost primary to much lower cost secondary storage tiers by policy, such as age or file size. Once in GridBank, a number of features ensure files are stored as efficiently as possible:

- Virtualization of underlying storage devices allows capacity on demand allocation of disk space. Dedicated capacity allocation results in unused headroom, which can occupy as much as 50% of total installed disk space.

- Single instance storage identifies and eliminates multiple copies of the same file. Data compression technology delivers further space saving.

- Application of lifecycle management policies places data in the right place, on the right media, at the right time based on its relevance and value to the organization.

- The simplified architecture reduces the number of locations across the organization where data is stored, reducing management overhead.

Benefits of Object Based Storage
GridBank stores unstructured file content such as Microsoft Office documents, PST files, images, and video surveillance files as objects. Each object consists of three main components: the file itself, metadata that describes the file, and configurable policies that manage numerous elements such as retention, security and access. Metadata provides an easy and fast method for the search and retrieval of data. Metadata can be created both for data held in the active archive, as well as files archived to tape. Since all metadata is held in the active archive, a single content search can traverse all archived files, both online or offline.
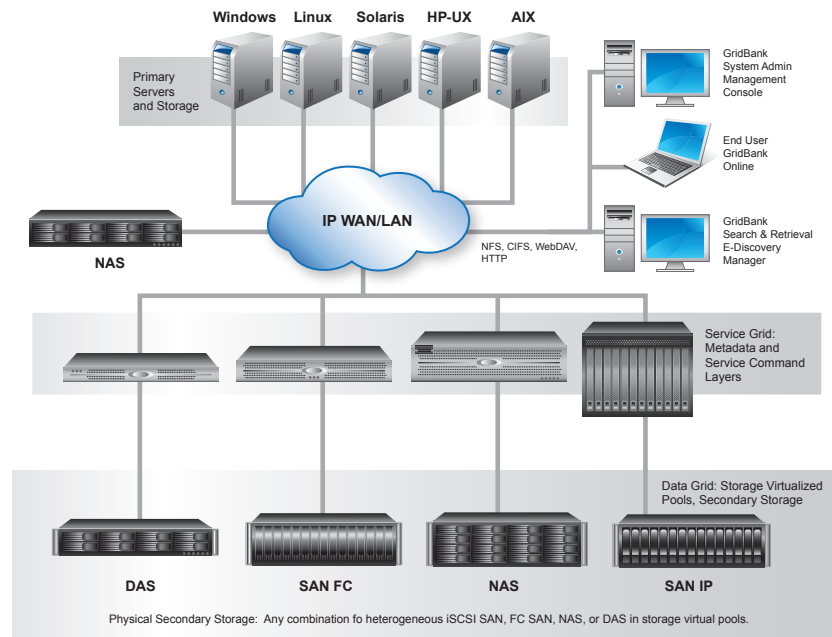
Eliminate Storage Inefficiency

GridBank automatically migrates data from expensive primary storage to more cost effective secondary storage. With primary storage costing up to five times as much as secondary storage, this can lower storage capital expenditures by as much as 80% per TB.   The volume of data that can be migrated is significant.   The Storage Networking Industry Association (SNIA) estimates that as much as 68% of primary storage is allocated to infrequently accessed data. GridBank's policy driven solution allows system administrators to set policies to automatically migrate data to the archive, based on criteria such as age, file size or type of file. Capacity on primary storage systems can be reclaimed for what it is intended for — high volume, high availability applications. Creating efficiencies in primary storage usage reduces the need to purchase additional capacity.   Reducing the volume of primary data substantially shortens backup windows, lowering the risk of backup failure. The result is better data protection and reduced risk of data loss. Once in GridBank, single instance storage, thin provisioning, and compression technologies deliver further savings by minimizing the secondary disk space required.

How it works

GridBank is a software only solution that is deployed on heterogeneous, industry standard server and storage components. It consists of two independent, yet interdependent elements: a service grid and a data grid. The service grid consists of a clustered group of servers, which hold the metadata layer and provide a pool of processing power to execute GridBank services. The data grid consists of underlying devices in a virtualized storage pool for easy allocation, provisioning, and management.

The data grid can include any combination of direct attached storage (DAS), storage area network (SAN), network attached storage (NAS) and tape or optical systems. Multiple device support allows GridBank to be deployed on any storage infrastructure and prevents vendor lock-in.

Figure 1: GridBank Architecture

GridBank Architecture: Physical View

Product Highlights

Cost Savings

- **Reclaim primary disk space**
- **Simplify administration**
- **Single instance storage**
- **Rapid data discovery**
- **Reduce storage footprint**
- **Environmental savings**

Product Highlights

Secure, Compliant Storage

- **Content authenticity**
- **Data protection**
- **Role based access**
- **Encryption**
- **Shredding**
- **Complete auditing**

Automated Information Lifecycle Management

Policy-based automation and management is an essential element of GridBank. With IT administrators struggling to manage storage costs effectively, GridBank provides a completely automated facility for
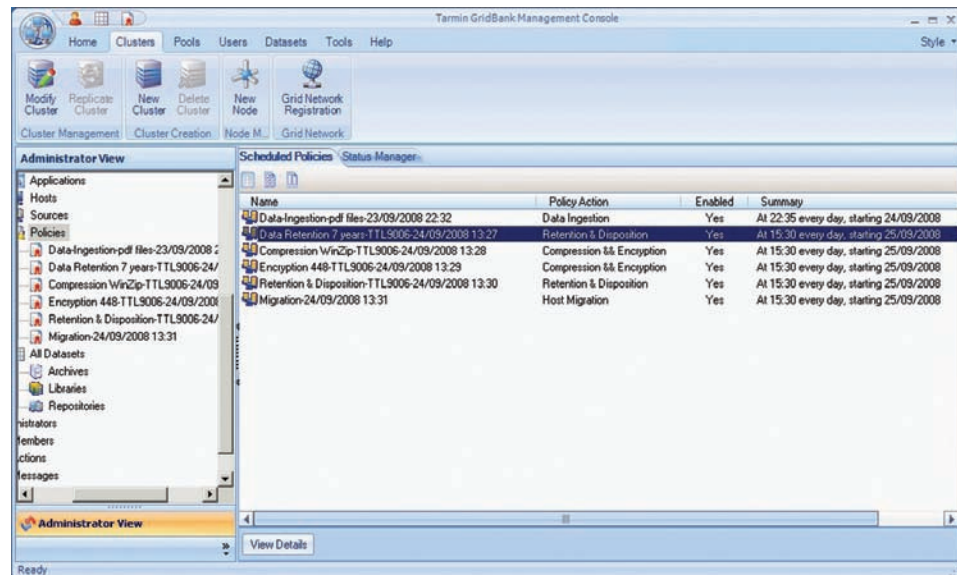
managing the lifecycle of file-based data. Once policies are set, all functions execute automatically, with no further storage administrator overhead. For example, a policy can be set to migrate files from primary storage to GridBank after 90 days, store in GridBank for one year, archive to tape for a further year, and then shred. Policies can even be scheduled to execute at non-peak times to minimize band-width utilization during peak network traffic periods.

Role Based Security, e-Discovery and Comprehensive Audit

GridBank's online client allows users to search content using an Internet style search engine. Searches are executed against the distributed metadata layer, which allows all GridBank data to be searched, wherever it is physically or geographically located. By adding files displayed in the search results to the user's shopping basket, copies are downloaded to the user's individual archive. If files are part of an ongoing legal investigation, a legal hold can be placed to prevent deletion. Users are authenticated via integration with corporate directory structures and data access policies are role-based, securing data from unauthorized access.

To insure the companies meet their legal and compliance requirements, GridBank includes a com-prehensive audit management facility. From the day data is migrated into GridBank, everything that happens to the archived object is tracked. Access, data movement from physical device to physical device, date of migration, and date of shredding are just some of the actions tracked in the GridBank audit manager.

Figure 2: GridBank Management Console

High Availability, Data Protection and Global Implementation

GridBank contains no single point of failure, providing high availability and ensuring data protection. Using a single global namespace, GridBank can be deployed enterprise wide, across multiple geographic locations. Support for industry standard protocols (NFS, CIFS) and a software development kit allows integration with third party applications as part of a strategic document management system.

" **The vast majority of business information, fixed content, is not being managed well, if at all. To extract value, control infrastructure costs and meet ever pressing government, regulatory and compliance mandates, information must be moved into a software managed active archive. GridBank's scalable, distributed robust architecture delivers the comprehensive function-ality required** "

— David Hill
Mesabi Group

FEATURE AND BENEFITS SUMMARY

Features	Benefits
Agentless Data Migration from Primary Storage	Save 75-80% per TB stored by migrating data from primary to secondary storage. Reducing primary volumes lowers risk of backup failure.
Storage Virtualization	Secondary storage devices are presented as a single or multiple virtual pools of storage for more efficient allocation of capacity and easier administration and management.
Clustered Grid	Load balancing architecture supporting the archive distributes services across all available compute resources in the grid, delivering high performance and allowing massive scalability (to 20 PB) while ensuring high availability.
Global Namespace	Allows geographic deployment of a single solution. Scales up to 20 PB and 32 billion objects.
Automated Policy Management	Reduces OPEX through automated policies throughout the data lifecycle, from creation to deletion to shredding.
Search Engine	Reduces the cost of locating data, increases productivity and reduces the risk of financial penalties for failing to respond to legal requests or audits within a specified time frame.
Single Integrated Solution	Single solution for an intelligently managed searchable archive and next generation secondary storage platform reduces cost and simplifies management compared to a multi-product approach.
Hardware Agnostic	GridBank runs on inexpensive commodity server hardware and can manage any combination of DAS, SAN, NAS, tape, and optical technologies, guaranteeing compatibility, lowering costs, and preventing vendor lock-in.
Single Instance Storage	Eliminates duplicate copies of files optimizing storage utilization.
Storage Efficiency	Compression of data delivers space savings. Thin provision reduces administration costs by cutting the time to provision new storage and allows for transparent addition of new storage capacity.
Compliance & Governance Support	Files are stored in a write once read many (WORM) format with a unique digital fingerprint that guarantees authenticity to meet compliance, governance, and regulatory requirements.
Data Retention Policy	Protects against premature deletion of data. Policies can be set to meet the requirements of individual organizations or business functions.
Role Based Access	Integration of GridBank with the existing corporate directory structures and identity management policies allows role-based, secure access to data.
Secure Storage	Data is encrypted to insure maximum data security.
Comprehensive Audit Management	Data authenticity is guaranteed through comprehensive auditing of files from data of ingestion to end of life.
Data Shredding	Data is properly destroyed at the end of its retention period, using digital shredding technology compliant with US Department of Defense data cleansing standards, providing protection against data leakage.
3rd Party Software Integration	Support for industry standard protocols (NFS, CIFS) and a software development kit allow integration with third party applications allowing full integration with other search and content management systems.

About Tarmin

Tarmin™ Technologies is a leading provider of active archival and next-generation intelligent storage software solutions optimized for secondary storage environments. Tarmin's GridBank utilizes industry standard, heterogeneous server and storage platforms to form a grid-based active archive and intelligent, scalable storage solution that enables effective long-term, fixed-content data preservation on cost effective secondary storage tiers. A high-performance, high-availability and highly scalable solution, GridBank satisfies any organization's regulatory compliance requirement by ensuring secure, long-term data retention, and fast search and retrieval of valuable business records. Tarmin is headquartered in Ongar, United Kingdom, with North American offices in Palo Alto, California. For more information, please visit **www.tarmin.com** or e-mail us at **info@tarmin.com**.

TARMIN

www.tarmin.com

United Kingdom
Essex Technology & Innovation Centre
The Gables
Fyfield Road
Ongar
CM5 0GA

Phone: +44 (0) 1277 368 333
Fax: +44 (0) 1277 368 205
Email: sales@tarmin.com
               info@tarmin.com

U.S.A.
228 Hamilton Avenue 3rd Floor
Palo Alto, CA 94301
United States

Phone: +1 (866) 4 TARMIN
Fax: +1 (650) 292-2223
Email: sales@tarmin.com
               info@tarmin.com