TARMIN TECHNOLOGIES ANNOUNCES GRIDBANK - ACTIVE ARCHIVAL AND INTELLIGENT STORAGE PLATFORM SOLUTION

Provides Increased Productivity, Compliance, Green Storage, and Reduces Overall Storage Costs

Ongar, United Kingdom and Palo Alto, Calif. (October 14, 2008) — Tarmin Technologies, a leading provider of active archive and intelligent storage software optimized for secondary storage environments, today announced GridBank™. This new and unique approach to enterprise-class active archiving and next-generation intelligent storage software allows organizations running Windows, Linux, Solaris®, HP-UX®, AIX®, or virtual servers (VMWare® or HyperV) to significantly reduce their storage Total Cost of Ownership (TCO).

With file-based data expected to be the largest percentage of all data stored by 2010, the high cost of primary storage, and the multitude of government regulations on data retention, the pressure on IT to lower budgets, control costs, and satisfy regulators is tremendous. GridBank solves these problems by dramatically reducing the cost and complexity of retaining, managing, and securely accessing unstructured data, while satisfying compliance requirements by ensuring secure, long-term data retention, e-discovery support, and fast search and retrieval of valuable business records. GridBank also delivers complete automation to increase productivity and substantially reduce storage and IT costs.

According to leading IT industry analyst firm Enterprise Strategy Group (ESG), digital archiving of data is growing at a 58% CAGR. With 70 percent of archived data being file-based unstructured data, ESG forecasts that digital data archiving will reach 90,000 petabytes (PB) in 2012, up from 9,197 PB in 2007 (Nov 2007).

“GridBank ™ provides a policy-driven, object-based archival solution that automates the movement of file-based data from expensive primary storage to much lower cost secondary storage,” said Brian Babineau, senior analyst at ESG. “GridBank also delivers the data retention, data preservation, audit tracking, security, and rapid content search and retrieval that lowers the cost of compliance and e- discovery. GridBank’s comprehensive integration of an active archival and intelligent storage platform means Tarmin can help solve customers’ archiving and information lifecycle management needs.”

GridBank Design and Architectural Overview

GridBank software utilizes industry standard, heterogeneous server and storage components to form a grid-based active archive and intelligent storage solution. Based on an object-based file model, GridBank incorporates automated, policy-driven information lifecycle management (ILM). With GridBank administrators easily create policies to automatically migrate data throughout its lifecycle - from one storage pool to another and from one type of storage media to another. GridBank’s highly fault tolerant solution intelligently scales up to 20 PB and over 32 billion individual objects using its globally distributed clustered file system. To substantially reduce storage costs, data is archived from expensive primary storage to much lower cost secondary storage.

With support for industry standard protocols, such as CIFS and NFS, GridBank provides an agentless design. Host server agents are not required on the primary servers and simplify deployment of GridBank. Data on

primary servers is migrated to GridBank over standard 1 Gb Ethernet infrastructure. GridBank software runs on Windows or Linux installed on physical or virtual servers and data can be automatically migrated to GridBank from any type of host storage - direct attached storage (DAS), storage area network (SAN), or network attached storage (NAS).

The GridBank platform stores data as unique objects (often referred to as Content Addressable Storage – CAS), which are indexed, and then administrator defined policies, such as encryption, role-based security access, data retention/preservation, single instance data deduplication, digital shredding, and compression, are automatically applied to the objects.

Grid-based Solution

GridBank's heterogeneous clustered server grid is networked together via 1 GB Ethernet. Servers in the grid can be from any supplier mixed-and-matched. GridBank's globally distributed clustered file system can be spread across multiple physical locations. Globally distributed archives provide both content distribution and archive disaster protection, while using existing secure WAN infrastructure.

A fault tolerant, self-healing design, GridBank archives continue to be available even in the event of hardware or software failure. Additionally, as demands on an archive grow, server infrastructure can be added non-disruptively and allow continued access while new resources are added. Server compute resources can be added to the existing grid or as a new cluster.

GridBank provides a unique load balancing architecture that optimizes server resources running GridBank's application services. In other load balancing architectures, system resources are matched to compute functions on a server-by-server basis – function 1 running on server 1 of the grid, function 2 running on server 2 of the grid and so on. With GridBank, however, services are distributed across all available computing resources so that processor power is fully optimized. If a particular GridBank application service needs the compute resources of two servers on one day but four servers on another day, the grid dynamically allocates compute resources to that application service as needed.

GridBank also supports multiple, simultaneous execution. A GridBank deployment can be concurrently ingesting data, searching the archive, compressing data, and non-disruptively adding new storage resources. This innovative design optimizes system resources, improves response time, and handles execution of simultaneous service functions better than other grid-based solutions.

Heterogeneous Storage Virtual Pooling

GridBank storage virtual pooling technology allows the creation of a single pool up to 20 PB or the creation of multiple pools. GridBank is fully media agnostic and employs heterogeneous storage pooling virtualization technology. Virtual storage pools can be created from any combination of SAN, DAS, NAS, tape, or optical subsystem from any manufacturer in a mix-and-match heterogeneous virtualized pool. This prevents vendor lock-in, giving IT maximum flexibility to buy whatever products meet its needs from any supplier. Storage resources can be added to and removed from virtual pools transparently while the archive is fully active.

GridBank delivers immutable data preservation and permanence through Write-once Read-many (WORM) design. This design is self-healing and ensures data is permanent. In the event of a physical storage failure,

data is replicated across other storage devices and a new replica is automatically generated. The secondary copy – known as a replica – allows GridBank to maintain WORM functionality and data preservation.

Storage Optimization, Efficiency and Green Storage Environment

GridBank optimizes storage, provides maximum efficiency of data, and help reduces carbon footprint, creating greener storage deployments. With primary storage subsystems costing up to four – five times more than secondary storage subsystems, CAPEX is instantly reduced with GridBank as data is moved from expensive primary storage to lower cost secondary storage. Since secondary storage requires less power and cooling than high performance primary storage and has a reduced carbon footprint, GridBank makes data center deployments much greener.

GridBank promotes CAPEX savings through its heterogeneous storage and virtual pooling design. Subsystems pricing will vary from vendor to vendor and over time as business conditions change. GridBank allows IT administrators to take advantage of these price changes and to combine storage from different vendors into the same virtual storage pool.

GridBank also delivers considerable OPEX benefits through its policy-based automation. As policies are executed, all storage functions in the GridBank solution are performed automatically, eliminating time consuming manual efforts and lowering IT manpower needs.

Further, GridBank delivers significant backup savings. According to a study conducted by the Storage Networking Industry Association (SNIA) 80 percent of all data is not changed after 90 days. Yet, by not moving data to an archive, that unchanging data would be backed up constantly, increasing backup windows, adding IT administrative costs, and slowing primary server performance. However, once data is archived to GridBank, it will no longer need to be backed up, saving time, money, and improving performance.

GridBank substantially reduces secondary storage utilization through the use of single instancing and data compression technology. GridBank archives only one instance of a file that actually has multiple copies on primary storage and then compresses that file. Multiple versions of the same file are eliminated, as is the actual size of that file. Secondary storage is significantly minimized and CAPEX significantly reduced.

Lastly, GridBank thin provisions storage used for its archives. As needed, physical storage can be non-disruptively added to GridBank's virtual pools. As additional capacity is needed, GridBank proactively alerts administrators that more physical capacity is required. This allows administrators to use a "pay as you go model" only buying additional storage resources as needed.

"Because GridBank offers key data management functions traditionally offered by separate point products, GridBank dramatically lowers the TCO of storage, helps organizations meet compliance needs, and makes the datacenter greener," said Shahbaz Ali, founder and CEO of Tarmin. "With the current turbulent economic conditions, organizations need highly efficient solutions, such as GridBank, to lower their CAPEX and OPEX, to increase IT manpower efficiency, and to secure their data with the latest technology for access and data sharing."

Policy-based Automation and Management

Policy-based automation and management are essential elements of GridBank. With IT administrators struggling to manage their storage cost effectively, GridBank provides a completely automated and highly customizable facility for managing the lifecycle of file-based data. Once the GridBank administrator sets policies, all functions are fully automated. Policy setting is so flexible in GridBank, policies can even be scheduled to execute at specific times. GridBank policies enable each archive to manage itself, independent of client applications and with no storage administrator overhead. GridBank supports the following policies:

- Data Protection - protects the integrity of data objects, initiating repairs immediately after any failure.
- Secure Data Access – with GridBank's identity centric design, all access is role-based and is integrated with existing Identity Management vehicles, such as Active Directory and LDAP.
- Data Authentication - ensures that the content of a file matches its digital signature and is valid data.
- Data Integrity - ensures that the content of a file is always good and accessible.
- Data Retention and Shredding – sets the length of time files are to be retained and prevents premature deletion of files. This policy allows for easy setting of retention based on legal requirements and corporate business practices. Data is shredded using the U.S. Department of Defense seven-pass shredding standard.
- Encryption – this standard feature supports encryption up to 448 bit. Encryption takes place at the object level to ensure the highest standard of security.
- Compression - reduces storage requirements and bandwidth usage.

E-discovery, Search and Audit

In order to optimize the business value of archived data and to leverage GridBank's active archival functionality to lower the costs associated with compliance, governance, and regulation, GridBank offers extensive searching, sharing and, e-discovery capabilities.

As part of the creation of objects in GridBank, contents of all files are fully indexed. This indexing, as well as the metadata associated with each object, allows GridBank to deliver a very fast search and retrieval capability. This is critical not only to lower the cost of e-discovery, but also to deliver superior value of the data that has been archived.

GridBank's online client allows users to search content using an Internet-style search engine. Searches are executed across archives wherever they are 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 ensure companies meet their legal and compliance requirements, GridBank includes a comprehensive audit management facility. From the day data is migrated with GridBank, everything that happens to the archived object is tracked, including access, data movement from physical device to physical device, date of migration, and date of shredding.

Summary

GridBank's next-generation active archive and intelligent storage software solution provides:

- Substantial reduction of IT, storage, and compliance CAPEX and OPEX
- All-in-one solution for CAS, ILM, active archiving, search and e-discovery, security, and audit management
- Increased productivity through automated policy-based archive and enhanced secondary storage management
- Fast search and retrieval for e-discovery
- Complete data lifecycle management from migration through to end-of-life
- Comprehensive storage optimization and efficiency to reduce the capacity of archived data, to deliver more efficient storage power and cooling, and to provide much greener storage environments
- Highly secure, yet easily accessible archives

Availability

GridBank will be generally available in early 2009 from Tarmin and its network of value channel partners.

ABOUT TARMIN TECHNOLOGIES

Tarmin™ Technologies is a leading provider of active archival and next-generation intelligent storage solutions optimized for secondary storage environments. Tarmin's GridBank™ platform solves archiving problems across all business functions. Deployed as a heterogeneous grid of clustered servers, GridBank delivers maximum data reliability and substantially lowers storage and IT CAPEX and OPEX. A high-performance, high-availability, highly scalable, and active archiving solution, GridBank satisfies any organization's regulatory, governance, and compliance requirements by ensuring secure, long-term data retention, and fast search and retrieval of valuable business records. GridBank's policy-driven software completely automates all processes, while empowering staff to securely share, search and publish their archived data. 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

Contact Information:

Corporate

Eric Herzog
Vice President of Sales and Marketing
Phone: +1 (510) 508-1871
Email: eric.herzog@tarmin.com
www.tarmin.com

North American Public Relations

Corey Oiesen, Julie Parayno
Dovetail Public Relations
Phone: +1 (408) 395-3600
Email: coreyo@dovetailpr.com, jp@dovetailpr.com

EMEA Public Relations

Federica Monsone
A_3 Communications
Phone: +44 (0) 1252 875 203
Email: fred.monsone@a3communications.co.uk