

Grid Backup for the Enterprise

Grid-X Backup is RAID's distributed grid storage platform that heals today's storage pains of infrastructure complexity, limited scalability, and management overhead through its "community of smart nodes."

A next-generation enterprise grid storage platform designed from the ground up, Grid-X Backup addresses key backup and archive challenges without the complexity, risks and operational limitations of current VTL, scale-out NAS, and other disk-based backup or archive appliances.

Grid-X Backup provides user configurable, integrated data management services to streamline storage management.

What makes Grid-X Backup unique is its distributed grid architecture. Combining revolutionary intelligent software, DynamicStor™ with best-in-class industry standard servers allows Grid-X Backup to deliver effortless management; unrestricted performance and capacity scalability; unmatched capacity optimization; and enhanced data resiliency.

Delivered as a turnkey solution, Grid-X Backup can be installed and backing up or archiving data in under minutes. Grid-X Backup is a fully self-managed, self-tuning and self-healing system that moves beyond "thin provisioning" to a new realm of "no provisioning."

Grid-X Backup's two-tier grid architecture delivers unrestricted, independent linear scalability of both performance and capacity based on business and application needs, and Grid-X Backup's self-evolving grid architecture eliminates legacy forklift upgrades and error-prone manual data migration by automatically recognizing new nodes and optimally distributing data among them. Grid-X Backup accomplishes this scalability by using two node types within the grid, and nodes can be added non-disruptively and without labor-intensive configuration or tuning.

In a single grid, Grid-X Backup can efficiently and securely store months' or even years' worth of backup and archive data. This unparalleled, "pay as you grow", cost-effective scalability is a benefit directly resulting from Grid-X Backup's grid architecture and its ability to incorporate either or both performance and capacity upgrades "on the fly".

Patent-pending data deduplication technology, DataRedux™ evolves data duplication to the next level. Combining sub-file, inline deduplication with additional compression, DataRedux achieves or greater space reduction for standard backups. More importantly, Grid-X Backup performs deduplication without the risk of data loss lurking as a hidden risk in competitive deduplication point products.

Grid-X Backup's patent-pending Distributed Resilient Data™ (DRD) technology protects deduplicated data beyond the limits of RAID technology. At the DRD default setting of three, your data is protected against three disk or node failures delivering percent more resiliency than RAID with percent less storage overhead while ensuring no performance degradation during data rebuild operations.

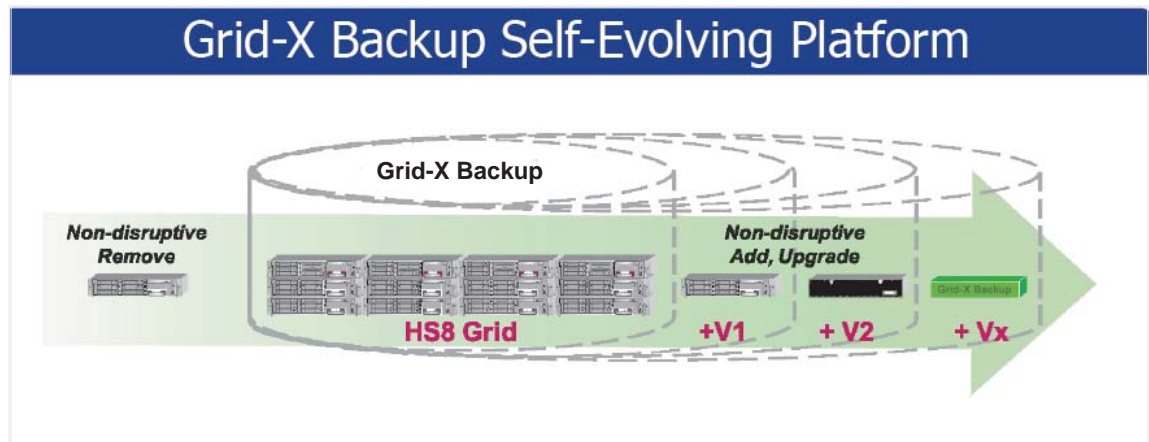
Another unique feature of Grid-X Backup is its ability to globally dedupe all incoming backup and archive data. By combining DataRedux with its grid storage architecture, Grid-X Backup delivers unified global deduplication across its entire grid, enabling to percent better dedupe efficiency than the siloed deduplication performance delivered by other deduplication products. Grid-X Backup is not only easy to manage, but easy to deploy, as it layers non-disruptively into almost any existing IT infrastructure.

As backup and archive data amounts grow, nodes can be non-disruptively added to increase search and retrieve performance as well as scale storage capacity. Each node is an industry-standard server with processing power built-in, so read/write performance is enhanced with each node addition as well.

With the HS-Platform Series, RAID delivers a new standard of storage offering possibilities previously thought impossible with current technologies. With its grid architecture and integrated data management services, Grid-X Backup offers a number of unique capabilities that reduce the cost and complexity of storage.

Through its use of industry standard access protocols, NFS and CIFS, Grid-X Backup integrates seamlessly with a wide variety of email, filesystem and document management applications and has been optimized to accept large volume data streams from backup and archive applications. Grid-X Backup has been validated with leading backup and archive software applications, including among others, Veritas NetBackup™, Symantec Backup Exec™, EMC® NetWorker®, Commvault Galaxy®, HP Data Protector, and Symantec Enterprise Vault™. Each Accelerator Node comes with multiple GbE or GbE NIC ports, conRAIDing Accelerator Nodes to one or more backup and/or archive servers to maximize throughput.

DynamicStor enables non-disruptive add, remove and upgrade functions and supports industry standards such as SNMP. DynamicStor eliminates many manual tasks, such as load balancing and capacity reservation, and offers transparent data



mobility. As nodes are added, capacity is automatically discovered, and existing data redistributed and load balanced to optimize performance and resiliency. Setup is quick and easy — Grid-X Backup can go from "box to backup and archive" in under minutes.


Grid-X Backup's community of smart nodes allows IT to non-disruptively ride the hardware technology curve, leveraging new node hardware without today's painful and time-consuming forklift upgrades, application downtime and manual data migration to new models. With Grid-X Backup, IT can introduce ever greener nodes, faster processors, and higher capacity hard disk drives into the same Grid-X Backup system, non-disruptively and painlessly in support of business needs.

Grid-X Backup's unique grid architecture enables both performance and disk capacity to be scaled independently to meet customer requirements. Accelerator Nodes scale performance from MB/s to MB/s; this performance coupled with data deduplication allows backup and archive windows to be reduced by to percent while enabling to percent faster access to stored data when needed. Storage Nodes scale storage capacity from TB to PB, with all storage managed as one logical pool of capacity, eliminating today's costly and labor-intensive siloed capacity challenges. With Grid-X Backup, IT has to manage only one system, instead of multiple systems with disparate interfaces, for both backup and archive data.

Grid-X Backup

Deduplication for Storage Efficiency

Grid-X Backup performs deduplication without the risk of data loss present in current point products. With the Distributed Resilient Data (DRD) technology “dial” set at its default level, Grid-X Backup provides protection against up to three concurrent HDD failures or an entire Storage Node with HDD. DRD further continues to maximize protection to up to three concurrent Storage Node failures (HDD) for configurations with or more Storage Nodes. This default setting – user configurable to provide higher or lower resiliency as desired – delivers times more resiliency than RAID with a similar storage capacity overhead, and times more resiliency than RAID with percent less storage capacity overhead while also ensuring virtually no performance degradation during data rebuild operations.



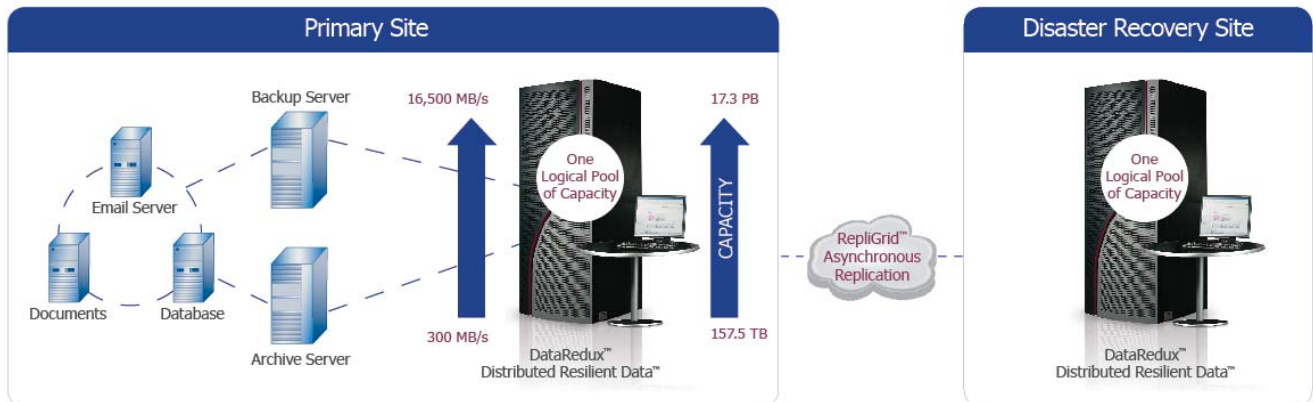
With the deployment of each additional appliance, competitive solutions with dedupe silos lose between 20 to 50 percent in both dedupe and capacity efficiency.

Grid-X Backup maximizes storage efficiency without the risk of data loss in current deduplication point products. Its patent-pending DataRedux technology eliminates redundant data across and within incoming data streams to enable Grid-X Backup to reduce storage requirements up to. Grid-X Backup's DataRedux also operates across both backup and archive data to eliminate redundancy between the two for further capacity savings.

RepliGrid is an optional feature to support disaster recovery requirements where data needs to be sent and stored at an offsite location. Grid-X Backup RepliGrid supports the ability to establish multiple independent grids that can share information with one another via asynchronous replication. Independent storage grids can be used to optimize local access for smaller remote offices and to create an independent disaster recovery (DR) facility for your primary data center.

In both cases, data is replicated to an alternate site to safeguard against natural disasters such as floods, fires, earthquakes, or storms that can affect the entire data center. Grid-X Backup RepliGrid provides an optimized cost-effective solution that enables efficient transfer of data by deploying bandwidth-friendly deduplication prior to replication over the WAN.

RAID's Grid-X Backup has enhanced replication capability with many to one replication functionality through Grid-X Backup RepliGrid™. Grid-wide deduplication minimizes network bandwidth requirements as well as disaster recovery site capacity requirements. With Grid-X Backup's granular data selection, customers can replicate just what they want, instead of being forced to replicate everything in their system at greater cost.



Grid-X Backup

Grid-to-Grid Replication for the Enterprise

Features

- Multi-site replication: enables data consolidation
- Duplicate-aware replication: only unique data chunks are sent across the network
- Data compression: reduces network bandwidth to replication sites
- Ability to schedule automated data resynchronization allows users to replicate data at the most opportune time

Benefits

Reduced Network Bandwidth Costs

- Grid-X Backup's duplicate-aware replication with data compression utilizes very low network bandwidth, enabling cost-effective disaster recovery

Reduced Storage Capacity Costs

- By consolidating multi-site data onto a single unified system with deduplication, storage capacity is reduced by 95 percent or more.

Simplified Network Management

- Grid-X Backup replication can be used to consolidate data from multiple sites for centralized processing.

Grid-X Backup's Unique Architecture

An enterprise grid storage platform architected, designed and optimized to address the challenges of both backup and archive. Grid-X Backup's unique grid architecture provides independent, unrestricted linear scalability of both capacity and performance while delivering global deduplication across the entire enterprise.

While Grid-X Backup features a new type of data resiliency called Distributed Resilient Data™ (DRD) technology to provide 300% more protection than RAID, RepliGrid is an optional feature to support disaster recovery requirements where data needs to be sent and stored at an offsite location. Grid-X Backup RepliGrid supports the ability to establish multiple independent grids that can share information with one another via asynchronous replication.

Why create independent storage grids?

1. to optimize local access for smaller remote offices
2. to create an independent disaster recovery (DR) facility for your primary datacenter

Grid-X Backup RepliGrid optimizes data transfer between grids by:

1. tracking changes since the previous synchronization point
2. verifying new data chunks do not already exist on the remote site
3. deduplicating data to ensure each unique chunk is only sent once
4. compressing the data to reduce information being transferred thus minimizing bandwidth requirements

Replicating information from a smaller remote office to a larger data center grid further improves the level of global deduplication by deduplicating information not only across file systems and nodes, but also across grids.

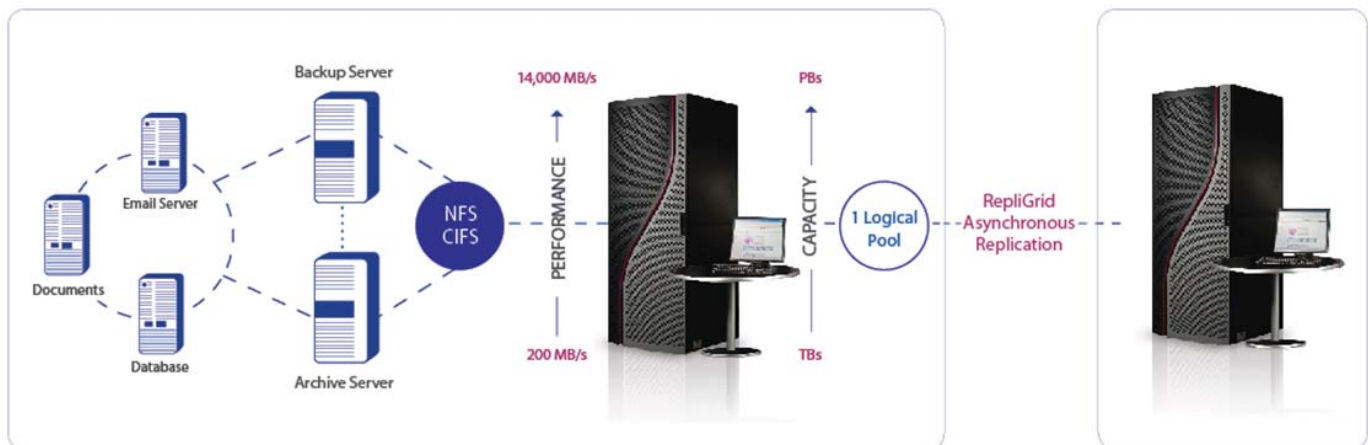
The replication process can be initiated manually or scheduled to take place at the most opportune time for your business cycle in order to ensure data availability and maximum performance while minimizing disruption to normal operations. Once new data is available at the remote site, Grid-X Backup's RepliGrid technology maintains a consistent, space-efficient copy of the data at all times. The remote site can also be used as a centralized facility for offloading data to tape or other removal media for offsite data vaulting.

Grid-X Backup is a self-evolving, self-managing, self-tuning and self-healing system that moves beyond "thin provisioning" to the new realm of "no provisioning." IT can non-disruptively ride the technology curve through the seamless addition of "ever greener", ever-faster, and ever-larger nodes to the Grid-X Backup grid.

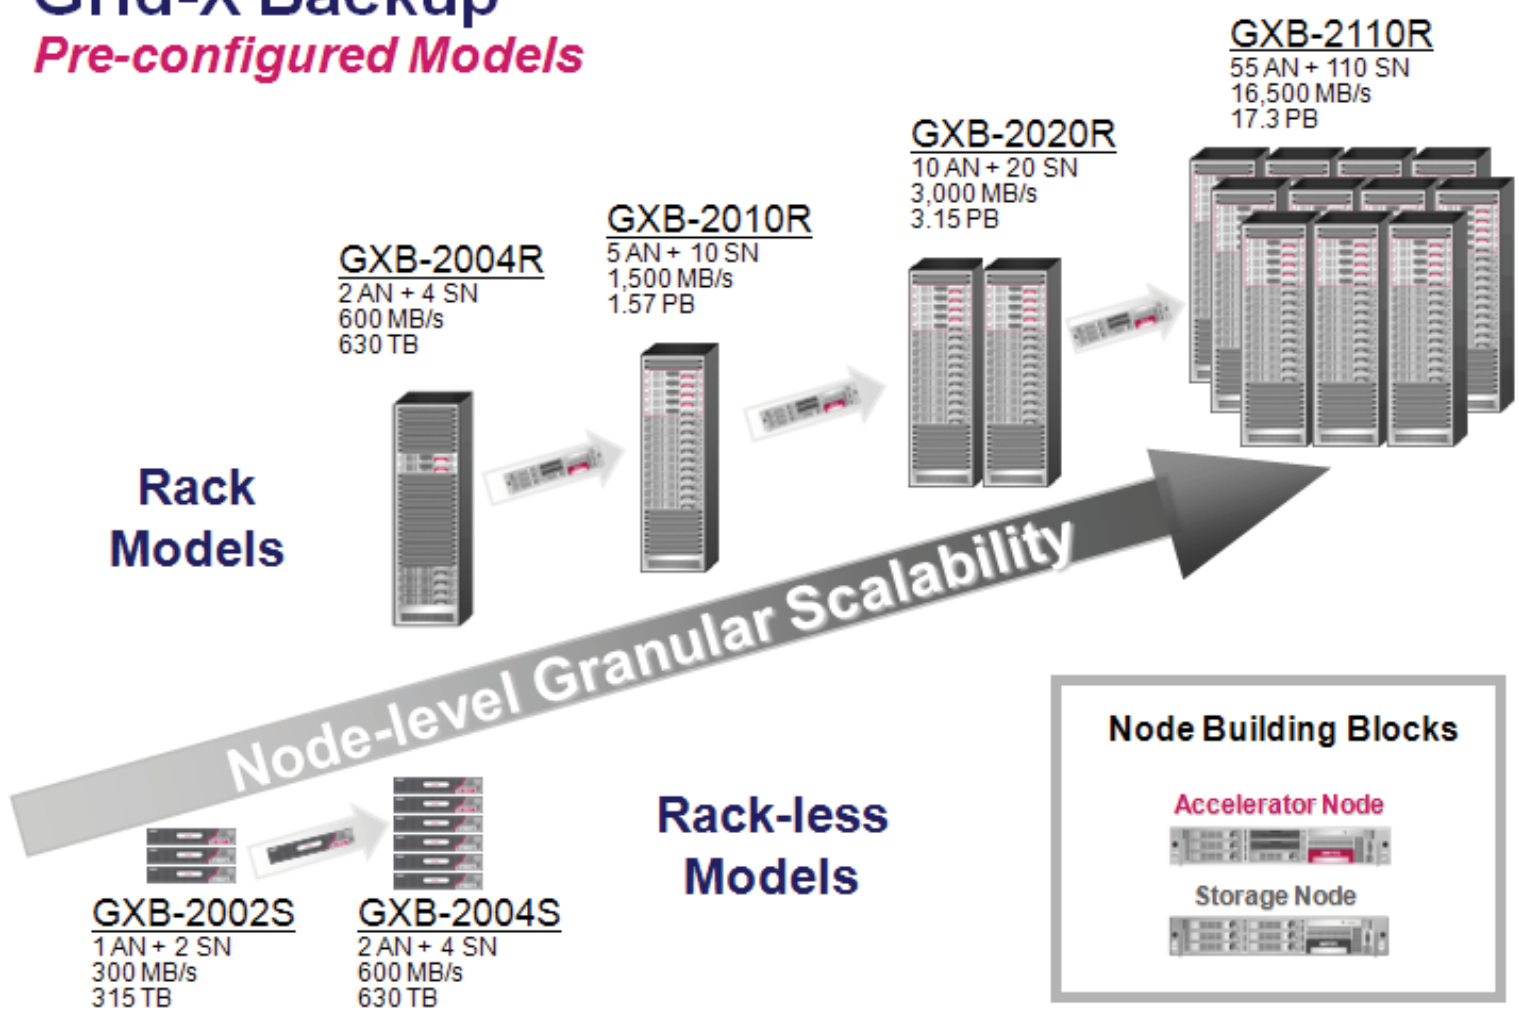
Grid-X Backup accomplishes unparalleled scalability by using two node types within the grid. Nodes can be added non-disruptively and without labor-intensive configuration or tuning.

Accelerator Nodes enable linear performance scalability for backup and archive

Storage Nodes allow linear capacity scalability with non-disruptive expansion from terabytes to petabytes. In a single distributed grid, Grid-X Backup can efficiently and securely store months' or even years' worth of backup and archive data. Affordable scalability with the ability to immediately leverage technology advances is a benefit of both Grid-X Backup's grid architecture and integrated data management services, including its patent-pending DataRedux™ and Distributed Resilient Data™ (DRD) technologies, incorporated into DynamicStor™, the software "brain" of Grid-X Backup™.



Grid-X Backup
Pre-configured Models



For more information on Grid-X Backup, call 800-330-7335 or email sales@raidinc.com