



DETAILED PRODUCT DESCRIPTION

ExaGrid Tiered Backup Storage

EXAGRID[®]
Tiered Backup Storage

Table of Contents

Executive Summary	2
ExaGrid Basic Concept	2
ExaGrid Technology Overview	3
ExaGrid Appliance Models.....	4
Scalable Architecture – Scale-Out.....	5
Designed for Performance.....	7
Data Deduplication	7
Data Backup Performance.....	7
Data Restore Performance.....	8
Support for Heterogeneous Backup Applications and Utilities	9
Advanced Backup Features.....	10
Veritas OST	10
Veeam Accelerated Data Mover.....	11
Veeam SOBR	12
Commvault	13
Oracle RMAN Channels.....	14
Reliability and Redundancy	15
ExaGrid Retention Time-Lock for Ransomware Recovery.....	16
RAID6 Internal Storage with Consistency Checking.....	19
Flash-Backed RAID Cache.....	19
Backup Data Checksums with Automatic Repair.....	19
Deduplicated Metadata Transactional Consistency.....	19
Internal Database Backups and Self-Describing Metadata	19
Logging Filesystem	19
Data Security	20
Active Directory Support	20
Periodic Assessments Using a Network Vulnerability Scanner.....	20
Offsite Data Protection for Disaster Recovery.....	21
Total Cost of Ownership.....	24
Cost Effectiveness Up Front.....	24
Cost Effectiveness as Data Grows	25
Summary.....	26
About ExaGrid	28

Executive Summary

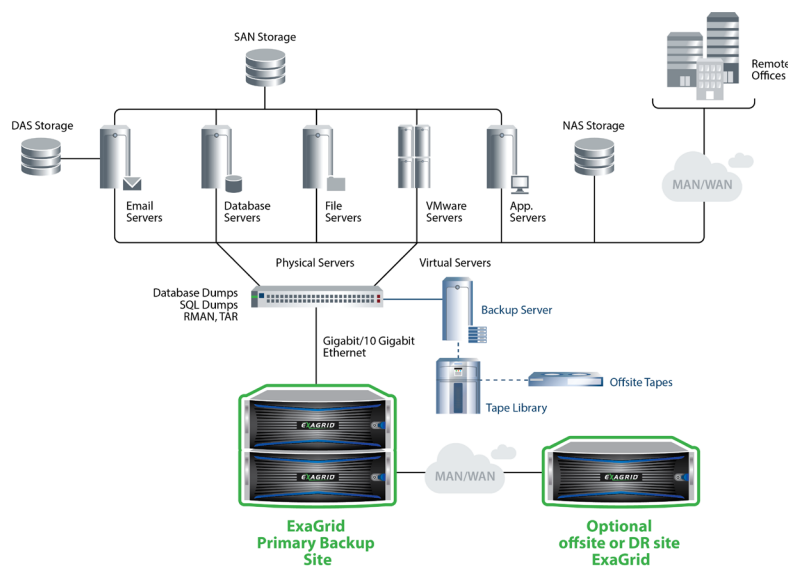
ExaGrid provides tiered backup storage with a unique disk-cache Landing Zone, long-term retention repository, and scale-out architecture. ExaGrid's Landing Zone provides for the fastest backups, restores, and instant VM recoveries. The disk-cache Landing Zone tier allows for the fastest backups and restores. The retention repository tier offers the lowest cost for long-term retention. ExaGrid's scale-out architecture includes full appliances and ensures a fixed-length backup window as data grows, eliminating expensive forklift upgrades and product obsolescence.

ExaGrid's Tiered Backup Storage makes backup better with:

- Fastest backups resulting in the shortest backup window
- Fastest restores
- Instant VM recoveries in seconds to minutes
- Fastest offsite tape copies
- A fixed length backup window as data grows, eliminating expensive future forklift upgrades
- Industry-leading deduplication ratios for storing long-term retention data for low-cost retention storage
- Deduplication efficiency across the entire system with global deduplication
- Fast, reliable, and up-to-date offsite disaster recovery using Adaptive Deduplication
- Linear scalability — scale-out architecture
- Low cost up front and over time as data grows
- No forklift upgrades or product obsolescence
- Retention Time-Lock to ensure recovery from a ransomware attack


ExaGrid Basic Concept

ExaGrid appliances work seamlessly with over 25 industry-leading backup applications and database utilities by presenting themselves as standard NAS shares (CIFS or NFS), a Veeam Data Mover target, or Veritas NetBackup OST target. ExaGrid appliances are easily integrated into existing backup environments, as illustrated below.



ExaGrid sits behind the existing backup server and replaces low-cost disk, inline deduplication appliances or tape backup storage, both onsite or offsite.

ExaGrid Technology Overview



ExaGrid Appliances Connect to Form a Single Scale-Out System

ExaGrid appliances are comprised of processors, memory, RAID6, and a hot spare drive per appliance, using enterprise-class SAS drives and ExaGrid software. See the **ExaGrid Technical Specifications** data sheet for more details.

Each appliance plugs in and is virtualized into a shared system with a single user interface, global deduplication, and automatic load balancing. The media server is connected to the same network and sees the appliances as one or more NAS shares, Veeam Data Mover targets, or Veritas NetBackup OST targets. Since each appliance includes the appropriate amount of processor, memory, Landing Zone disk, deduplicated repository disk, and bandwidth for the rated data size, performance increases as more appliances are added to the system.

ExaGrid Appliance Models

The product line's multiple appliance models can be combined into a single scale-out system with a configuration of up to 32 appliances in a single system for a total of 6.14PB raw capacity, 5.37PB of usable capacity, and supports a full backup of up to 2.69PB, total backup logical data capacity of 43PB.

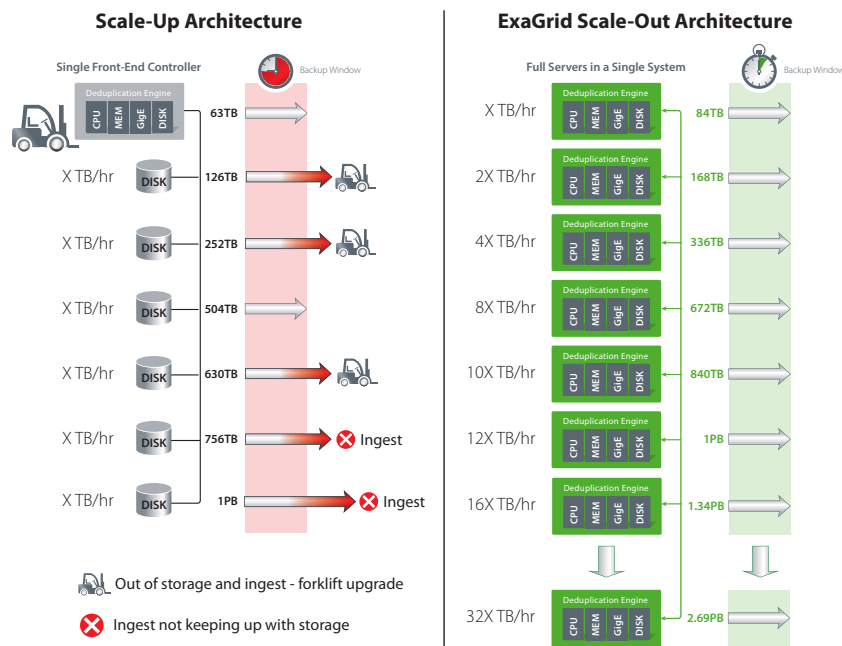
ExaGrid Model	Raw Capacity (TB)	Usable Capacity (TB)	Capacity for Weekly Full (TB)	Capacity for 16 Weekly Fulls (TB)	Max Backup Thruput (TB/hr)
ExaGrid Appliances					
EX6	24	12	6	96	1.74
EX10	32	20	10	160	2.40
EX18	48	36	18	288	3.60
EX27	72	54	27	432	6.09
EX36	96	72	36	576	7.78
EX52	128	104	52	832	10.87
EX84	192	168	84	1344	15.25
ExaGrid Appliances with Encryption					
EX10-SEC	32	20	10	160	2.40
EX18-SEC	48	36	18	288	3.60
EX27-SEC	72	54	27	432	6.09
EX36-SEC	96	72	36	576	7.78
EX52-SEC	128	104	52	832	10.87
EX84-SEC	192	168	84	1344	15.25
	TB / PB	TB / PB	TB / PB	PB	TB/hr
Example System Configurations					
EX168-G	384TB	336TB	168TB	2.68	30.5
EX504-G	1.15PB	1PB	504TB	8.06	91.5
EX1008-G	2.3PB	2.01PB	1PB	16.12	183
EX1344-G	3.07PB	2.68PB	1.34PB	21.5	244
EX1680-G	3.84PB	3.36PB	1.68PB	26.88	305
EX2100-G	4.8PB	4.2PB	2.1PB	33.6	381.25
EX2688-G	6.14PB	5.37PB	2.69PB	43	488

*Scales to 32 appliances in a single scale-out system
(2.69PB full backup @ 488TB/hr. ingest rate)*

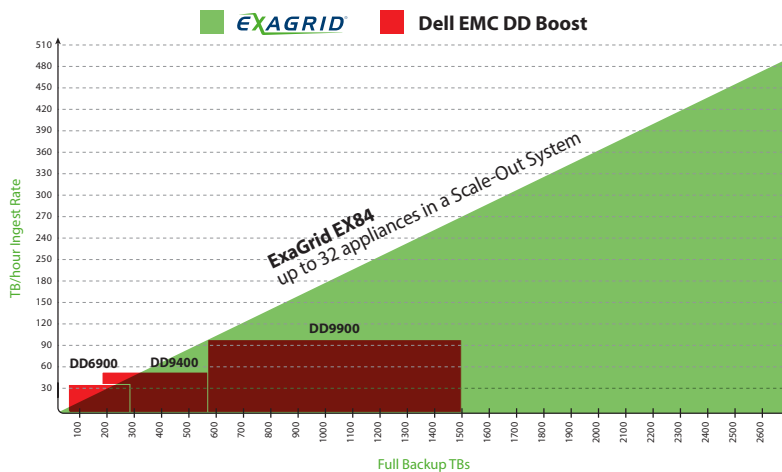
Scalable Architecture – Scale-Out

The typical business or organization is seeing data growth of 30% a year which doubles data every 2.5 years. This data growth presents challenges to IT to ensure their backup system can scale easily to support that growth. With scale-up disk backup solutions that have a front-end server/controller architecture and add just disk shelves as data grows, you start with a short backup window, but as data grows, the backup window expands because only disk is added but no additional deduplication processing resources. Eventually, the backup window expands to a point where you must replace the front-end server with a more powerful server via a costly “forklift upgrade.”

In contrast, ExaGrid’s scalable approach with global deduplication adds full servers—including memory, processor, and bandwidth as well as disk. The figure below shows the differences between how the two different architectures cope with data growth over time.

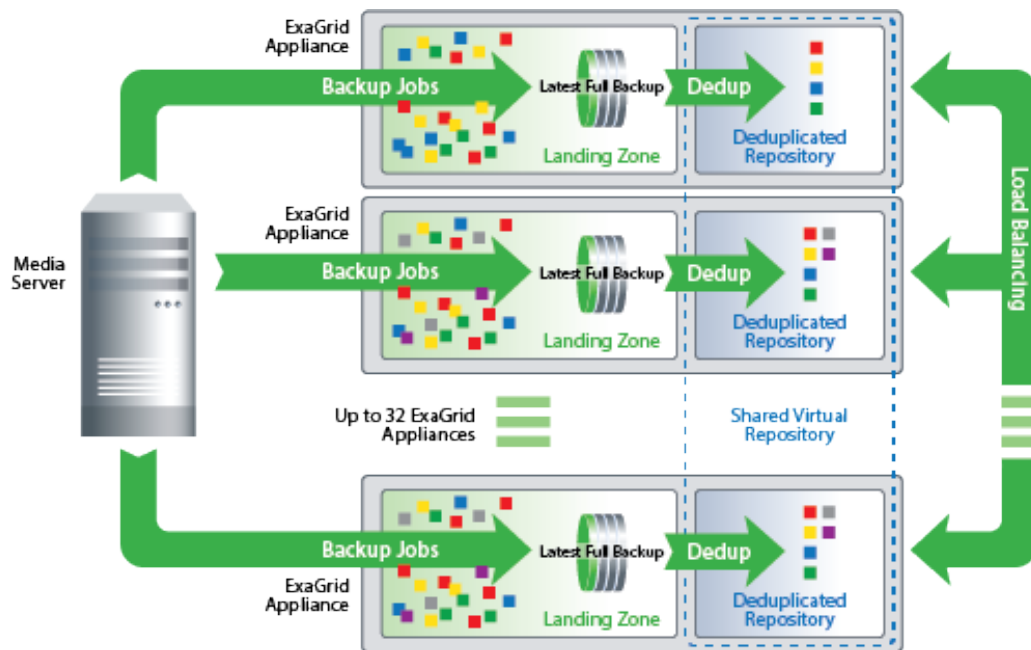


Scale-out vs Scale-up Backup Storage



This unique approach provides the following benefits:

- **No Expansion of Backup Windows as Data Grows** – By adding full appliances, ExaGrid maintains consistently fast backup performance and a fixed length backup window as data increases.
- **Seamless Scalability with No Forklift Upgrades** – ExaGrid’s scale-out approach avoids forklift upgrades by adding modular capacity with full servers to the existing system in full backup increments of 6, 10, 18, 27, 36, 52, or 84TB with up to 32 appliances in a single system that scales to 2.69PB full backup.
- **No Obsolescence of Previous Model Systems** – All ExaGrid appliances can be seamlessly added to any existing ExaGrid deployment and can be mixed and matched with older appliances of any size, thereby eliminating obsolescence. Adding a new ExaGrid appliance is as simple as plugging it in and pointing backup jobs at the appropriate target. If any model is no longer offered, the models still work in a scale-out system with older and newer models, maintenance and support is always offered and maintenance and support will not go up more than 3% per year allowing organizations to get their full lifetime use out of all appliances in the system.



ExaGrid appliances include scalable computing software, which allows them to virtualize and share data storage capacity with one another (automatic load balancing across all repositories). This scalable system (shown above) can expand as data grows by adding appliances, providing up to 6.14PB raw capacity, 5.37PB usable capacity, and allowing full backups of up to 2.69PB in a single scale-out system. Multiple systems can be used at a single location, and up to 16 separate locations can be managed through a single user interface.

Designed for Performance

Data Deduplication

ExaGrid's innovative approach to data deduplication minimizes the amount of data to be stored by using zone-level data deduplication across all received backups.

ExaGrid's zone-level technology stores only the changed data at a granular level from backup to backup instead of storing full copies. ExaGrid uses zone stamps and similarity detection.

This unique approach reduces the disk space required by an average of 20:1 and from 10:1 up to 50:1 depending on data type, retention and backup rotation delivering unparalleled performance for the fastest backups and restores.

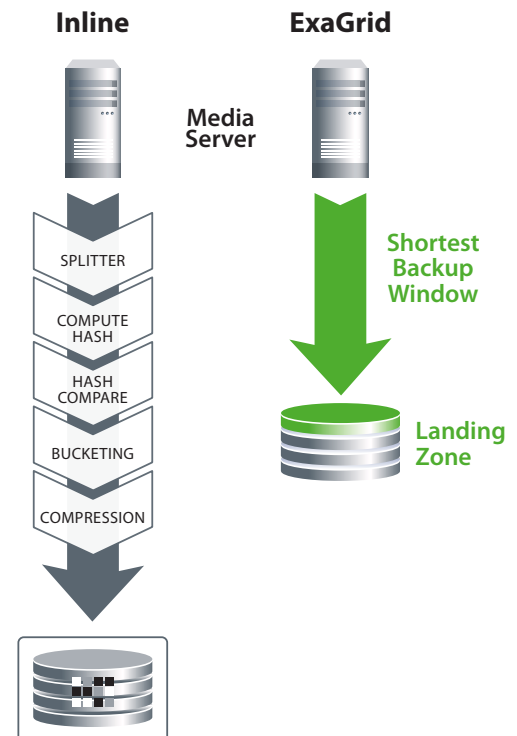
ExaGrid is the only vendor that allows backup applications such as Veeam and Commvault to keep their deduplication turned on and then ExaGrid further deduplicates the data for increased storage efficiency.

Data Backup Performance

The ExaGrid system performs deduplication in parallel with backups by using idle system cycles. Termed "adaptive deduplication," this approach keeps deduplication out of the backup path, resulting in faster backups and a shorter backup window. Adaptive Deduplication gives full system resources to backups for the shortest backup window and uses idle cycles to perform deduplication and offsite replication, allowing data to be replicated offsite for a strong offsite recovery point (RPO) for update to data for disaster recoveries.

ExaGrid supports data deduplication globally across all targets and all appliances in a scale-out system. Global deduplication ensures that all data is deduplicated regardless of the number of targets and/or the number of appliances in a system. In addition, global deduplication allows organizations the flexibility to redirect backup jobs to any target, on any appliance, at any time while maintaining data deduplication globally across the entire system.

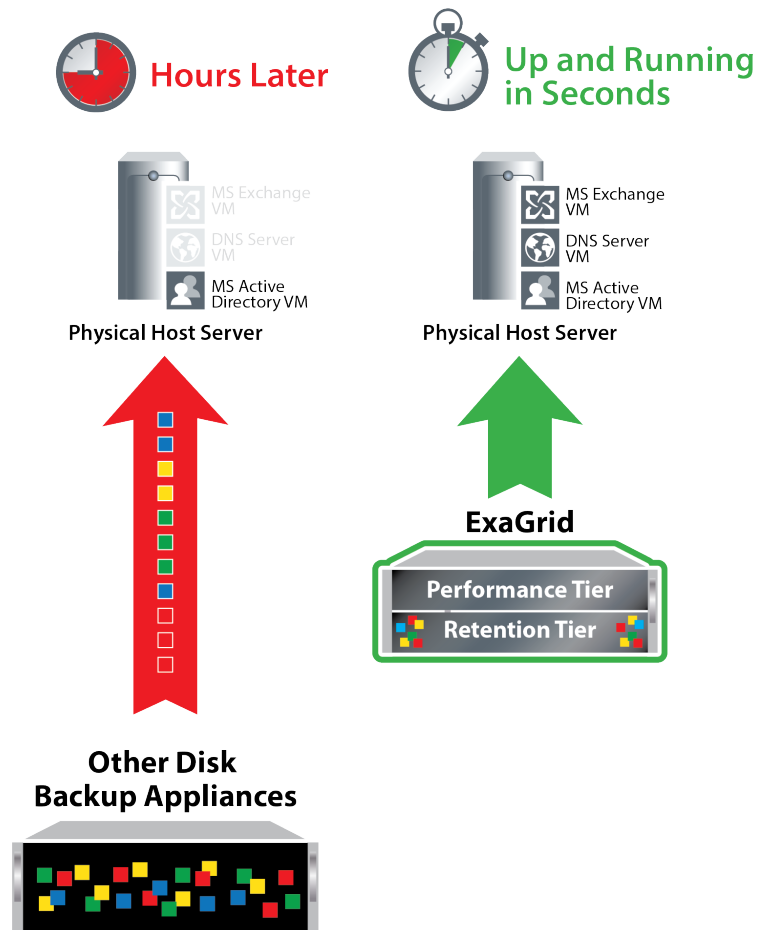
Once the backup job is complete and off the network, the data is protected and immediately available to be restored. When deduplication and replication are complete, the data is ready for restore on the ExaGrid appliances at the disaster recovery site. ExaGrid uses Adaptive Deduplication that avoids the performance penalty of inline deduplication. ExaGrid deduplicates in parallel with backups coming in but performs no inline processing which allows for at least a 3X performance advantage over inline deduplication appliances. As shown in the diagram to the right, backup data is written directly from the backup server to ExaGrid's disk-cache Landing Zone at the highest possible rate with no inline processing to interfere, resulting in the shortest possible backup window.



Data Restore Performance

Because ExaGrid's appliance allows each full backup to first land on the disk-cache Landing Zone before deduplication, the system maintains the most recent backup in its full, undeduplicated form for fast restores, Instant VM recoveries in seconds to minutes, and fast offsite tape copies. Since over 90% of restores and 100% of instant VM recoveries and tape copies are done from the most recent backup, this approach avoids the overhead incurred from "rehydrating" data during critical restores. As a result, restore, recovery, and copy times from an ExaGrid system are an order of magnitude faster than solutions that only store deduplicated data.

Restore Performance




Support for Heterogeneous Backup Applications and Utilities

Customer environments are made up of many backup approaches, backup applications, and utilities and different disk-based backup approaches support these in different ways. Customers may have any number of backups occurring in their environment, including traditional backup applications such as Veeam, Veritas NetBackup, Commvault, IBM Spectrum Protect (TSM), Dell EMC NetWorker, HYCU, Zerto, and many others as well as direct-to-disk SQL dumps, Oracle RMAN direct backups, and specific UNIX utilities such as UNIX TAR. ExaGrid supports over 25 backup applications and utilities.

ExaGrid is able to support backup data from multiple sources, including a variety of backup applications and database dump utilities. Performing deduplication in the backup software limits the ability to have all data from all sources stored and deduplicated in a single target device. Unless 100% of your backup data passes through that particular backup application, a purpose-built disk-based backup appliance such as ExaGrid's is the best choice to meet the requirements of your entire environment.

In contrast, backup application software solutions that have incorporated deduplication by definition only support their own backup application, with its own backup server software and its own backup client agents. These solutions are not able to support backup data from other backup applications or utilities. For example, if you have a physical environment, for backup applications that employ data deduplication but want to use a separate utility for VMware (ex: Veeam) and also do direct database dumps, only the data running through the physical system's backup application will be deduplicated. Also, deduplication in backup software ranges from 2:1 to 10:1 and therefore uses a lot more disk and bandwidth as retention grows versus target-side deduplication appliances that employ far more aggressive deduplication algorithms, since with appliances, resources are dedicated to the task. Dedicated appliances achieve ratios on average of 20:1, using one-quarter to one-third the disk and bandwidth.

ExaGrid allows for Veeam and Commvault to leave deduplication turned on and ExaGrid will further deduplicate the data for increase storage efficiency.



Advanced Backup Features

Veritas Open Storage Technology (OST)

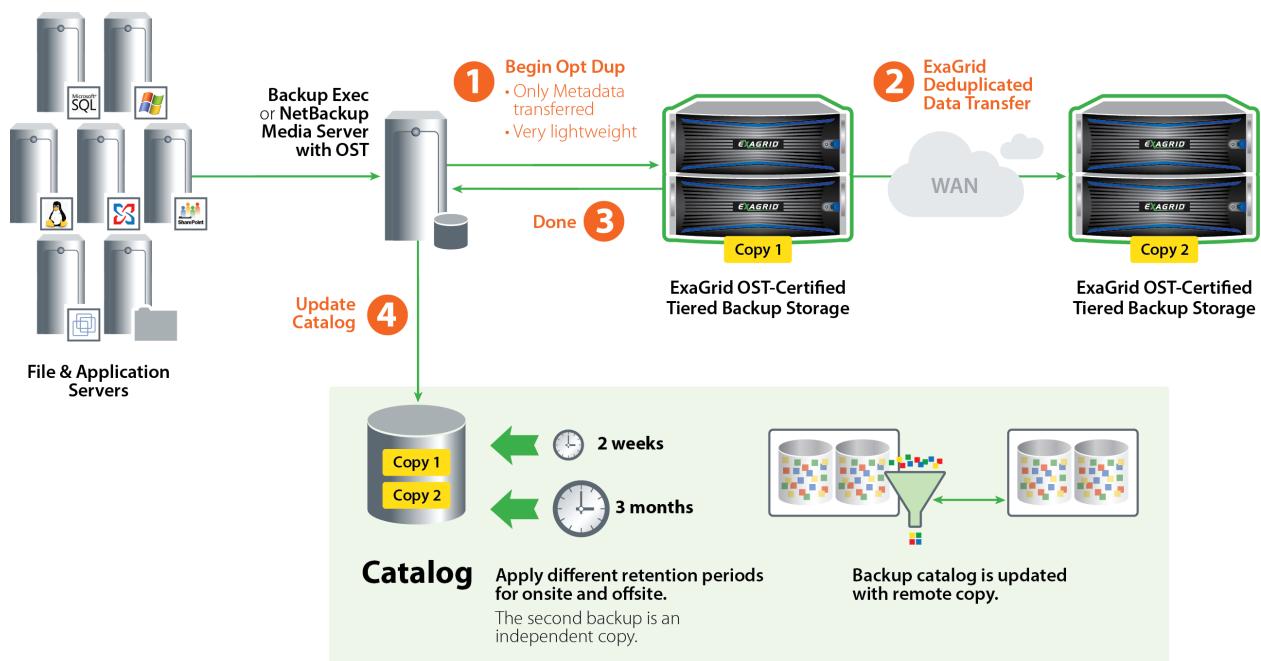
An important consideration when looking at disk-based backup solutions is how well a particular solution supports advanced backup application features such as Backup Exec GRT (Granular Restore) and Veritas' OST (Open Storage Technology) for Backup Exec and NetBackup. Some solutions do not integrate well with these features; poorly implemented GRT solutions, for example, may take hours to restore an individual e-mail or may not work at all.

Veritas' Open Storage Technology is another popular feature that allows for more integrated offsite data protection, and it is important to check whether these features are supported if you are using Veritas NetBackup or Backup Exec, OST allows for faster performance, better management, and unbalanced onsite and offsite retention.

ExaGrid supports and is certified for:


- NetBackup and Backup Exec OST
- NetBackup AIR
- NetBackup appliance compatibility
- NetBackup Accelerator

ExaGrid is the only solution in the market for Veritas NetBackup Accelerator that can reconstitute a full backup in its Landing Zone so that a complete backup is ready for restore in its already hydrated form for faster restores and VM boots.




Veeam Accelerated Data Mover

Most of Veeam's unique features such as its Scale-Out Backup Repository (SOBR), Sure Backup, Virtual Lab, Instant VM Recovery, Copy and Replicate, and other advanced features require an unduplicated copy on disk. Only ExaGrid provides this with its unique disk-cache Landing Zone. All other solutions only store deduplicated data. In addition, ExaGrid includes an integrated Veeam data mover with each appliance called the "ExaGrid-Veeam Accelerated Data Mover." This improves all backup and restore processes, increase security with a closed end-to-end communications protocol and also allows a synthetic full to be created directly on the ExaGrid system for increased performance.



Veeam Scale-Out Backup Repository (SOBR)



Veeam's Scale-Out Backup Repository (SOBR) allows backup administrators using Veeam to direct all jobs to a single repository made up of ExaGrid shares across multiple ExaGrid appliances with global deduplication in a scale-out system, automating job management to ExaGrid appliances. ExaGrid's support of SOBR also automates the addition of appliances into an ExaGrid system as data grows by simply adding appliances to a Veeam repository group. The combination of Veeam SOBR and ExaGrid's appliances in a scale-out system creates a tightly integrated end-to-end backup solution that allows backup administrators to leverage the advantages of scale-out in both the backup application as well as the backup storage. The combination of Veeam backups to the ExaGrid disk-cache Landing Zone, the integrated ExaGrid-Veeam Accelerated Data Mover, and ExaGrid's support of Veeam SOBR is the most tightly integrated solution on the market for a scale-out backup application to scale-out backup storage.


Commvault

ExaGrid allows for Commvault deduplication to be turned on. ExaGrid further deduplicates the Commvault data by a factor of 3x greatly reducing the amount of backup storage required.

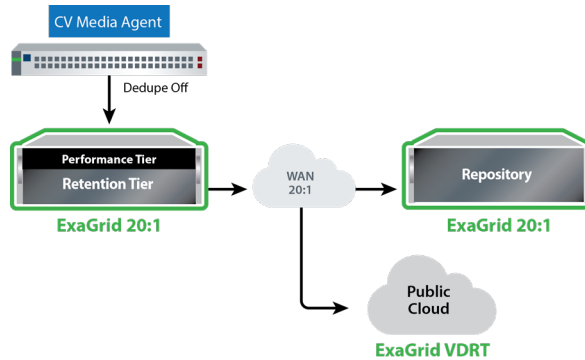
ExaGrid can also allow users to turn Commvault deduplication off to increase backup performance while retaining the same cost storage as leaving Commvault deduplication on with ExaGrid's additional deduplication impact.

ExaGrid support Commvault Spill & Fill for automatic job management where all jobs are sent to ExaGrid appliances in the system by Commvault automatically. Jobs can be sent to any appliance at any time as ExaGrid has both global deduplication across all appliances in the system and automatic load balancing of all long term retention data repositories.

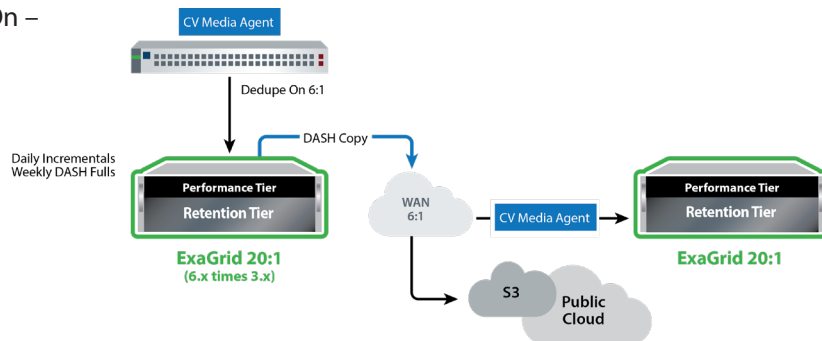
Commvault – Dedupe On



Commvault – Dedupe Off



Commvault Dedupe On – DASH Copy




Oracle RMAN Channels

ExaGrid supports Oracle RMAN Channels targeted at multiple NAS shares across multiple appliances with global deduplication in a scale-out system. RMAN Channels automatically writes “sections” in parallel to all NAS shares and automatically redirects the next “section” based on available targets. RMAN Channels with ExaGrid has six major advantages.

1. Oracle database can be up to 2.69PB in size and can be backed up in parallel to a single ExaGrid scale-out system.
2. The database backup performance is accelerated as the sections are backed up in parallel across multiple appliances in a scale-out system.
3. The database backup performance is maximized as each new section is automatically sent to the highest performance availability NAS shares and/or appliance, resulting in the best possible performance based on NAS share and appliance ingest availability.
4. If any appliance fails, the segments are automatically redirected to the active appliance, providing for automatic failover.
5. The most recent database is stored in an undeduplicated form in the ExaGrid disk-cache Landing Zone, allowing for fast restores while still allowing for storage efficiency as all long-term retention data is stored in deduplicated form. This avoids the lengthy data rehydration process of inline scale-up appliances that only store deduplicated data.
6. As the database data grows, the backup window stays fixed in length as full appliances are added into a scale-out system bringing compute with capacity. This eliminates the forklift upgrades associated with inline scale-up deduplication appliances.

Database Backup Performance



Reliability and Redundancy

Organizations using a disk-based backup appliance to hold their invaluable backup data should carefully consider how the appliance is architected for reliability and redundancy. Compromises in a product's architecture or implementation may reduce product cost, but those savings are quickly negated by the risk and real cost to an organization of a loss of some or all backup data.

ExaGrid's architecture and implementation have multiple facets of reliability and redundancy, allowing organizations that are considering disk-based backup appliances to make informed vendor selections.

ExaGrid offers the following ease of use, redundancy and security features, some of which are explained below:

- Single user interface for all appliances in a system and across sites
- RAID6 protection with a hot swappable spare
- Redundant hot swappable power supplies
- Active Directory for management interface and backup target security
- SNMP and syslogging interface for integration with enterprise management apps
- Role-based access control
- Retention Time-Lock – ransomware recovery
- Two-factor authentication
- Data encrypted at rest
- Data encryption while replicating over the WAN
- Security checklist makes it easy to apply best practices
- Data is checksummed to ensure data integrity
- Internal self-describing database

ExaGrid Retention Time-Lock for Ransomware Recovery

ExaGrid's unique approach ransomware recovery is called Retention Time-Lock. It prevents hackers from deleting the backups and allows for retention points to be purged. The result is a strong data protection and recovery solution at a very low cost of storage.

ExaGrid provides Tiered Backup Storage with a front-end disk-cache Landing Zone and separate Retention Tier containing all retention data. Data is written directly to the "network facing" ExaGrid disk-cache Landing Zone. Then it is tiered into a "non-network facing" long-term retention repository where it is stored as deduplicated data objects to reduce the storage cost of long-term retention data. As data is tiered to the Retention Tier, it is deduplicated and stored in a series of objects and metadata. As with other object storage systems, the ExaGrid objects and metadata never change allowing only for the creation of new objects or deletion of old objects when retention is reached.

ExaGrid's approach to ransomware allows organizations to set up a time-lock period that governs the processing of any delete requests in the Retention Tier as that tier is not network facing and not accessible to hackers. The combination of a non-network facing tier, a delayed deletion for a period of time and objects that never change are the elements of the ExaGrid Retention Time-Lock solution. For example, if the time lock period for the Retention Tier is set to 10 days, then when delete requests are sent to the ExaGrid from a backup application that has been compromised or from a hacked CIFS or other communications protocols, the data in the Retention Tier is time-locked for up to 10 days against any deletion. The data in the Landing Zone will be deleted or encrypted, however, the Retention Tier data is not deleted upon an external request for the configured period of time. When a ransomware attack is identified, simply put the ExaGrid system into a new recover mode and then restore any and all backup data to primary storage. The time lock period is separate and in addition to the days, week, months and year or retention that is set by the backup application and stored by ExaGrid in the retention repository.

The solution provides a retention lock, but only for an adjustable period of time as it delays the deletes. ExaGrid chose not to implement Retention Time-Lock forever because the cost of the storage would be unmanageable. ExaGrid already has the long term backup retention so it would be redundant to have a separate store with retention lock. With the ExaGrid delayed delete approach, all that is needed is up to an additional 6% more repository storage to hold the delay for the deletes. ExaGrid allows the delay of deletes from 1 day to 30 days.


ExaGrid advantages are:

- Manage a single system instead of multiple systems for both backup storage and ransomware recovery
- Unique second Retention Tier that is only visible to ExaGrid software, not to the network
- Data is not deleted as delete requests are delayed and therefore ready to recover after a ransomware attack
- Weekly, monthly, yearly and other purges still occur to keep storage costs in line with the retention periods
- Only requires up to an additional 6% of repository storage
- Storage does not grow forever and stays within the backup retention period set to keep storage costs down
- All retention data is preserved and is not deleted

Example Scenarios

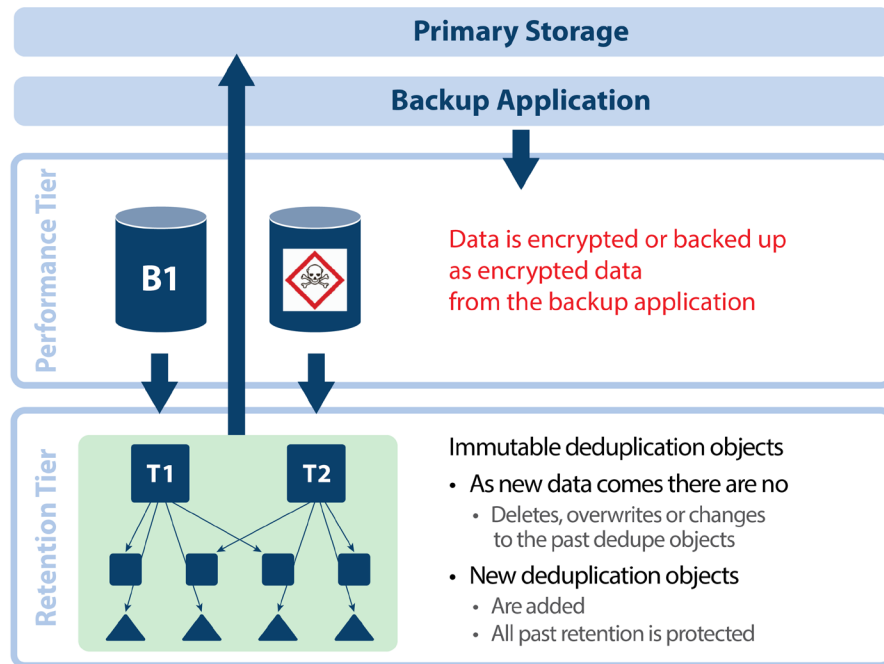
- 1) Data is deleted in the ExaGrid disk-cache Landing Zone via the backup application or by hacking the communication protocol. Since the Retention Tier data has a delayed delete time lock, the objects are still intact and available to restore. When the ransomware event is detected, simply put the ExaGrid in a new recover mode and restore. You have as much time to detect the ransomware attack as the time lock was set for on the ExaGrid. If you had the time lock set for 10 days, then you have 10 days to detect the ransomware attack and put the ExaGrid system in the new recover mode for restoring data.

Deletion Protection of Backup Data on ExaGrid



- 2) Data is encrypted in the ExaGrid disk-cache Landing Zone or is encrypted on the primary storage and backed up to ExaGrid such that ExaGrid has encrypted data in the Landing Zone and deduplicates it into the Retention Tier. The data in the Landing Zone is encrypted. However, all previously deduplicated data objects never change (immutable), so they are never impacted by the newly arrived encrypted data. ExaGrid has all previous backups before the ransomware attack that can be restored immediately. In addition to being able to recover from the most recent deduplicated backup, the system still retains all the backup data according to the retention requirements.

Deletion Protection of Backup Data on ExaGrid



Features:

- Any deletion requests are delayed by the number of days in the protection policy.
- Encrypted data written to ExaGrid does not delete or change previous backups in the repository.
- Landing Zone data that is encrypted does not delete or change previous backups in the repository.
- Set delayed deletion in 1 day increments from 0 days to 30 days.
- Protects against loss of any and all retained backups including monthlies and yearlies.
- Two-Factor Authentication (2FA) protects changes to Time-Lock setting.
 - Only Security Officer role is allowed to approve changes to Time-Lock setting.
 - 2FA with Login/Password and system generated QR code protects all accounts.
- Separate password for primary site versus second site ExaGrid.

Special Feature: Alarm on Delete

- An alarm is raised 24 hours after a large delete.
- Alarm on large delete: A value can be set as a threshold by the backup administrator (default is 50%) and if a delete is more than the threshold, system will raise an alarm, only Admin role can clear this alarm.
- A threshold can be configured, by individual share, based on backup pattern. (The default value is 50% for every share). When a delete request comes to the system, the ExaGrid system will honor the request and delete the data. If RTL is enabled, the data will be retained for the RTL policy (for the number of days set by an organization). When RTL is enabled, organizations will be able to recover the data using the PITR (Point-In-Time-Recovery).
- If an organization gets false positive alarm frequently, the Admin role can adjust the threshold value from 1-99% to avoid more false alarms.

RAID6 Internal Storage with Consistency Checking

All ExaGrid internal storage is accessed using an industry-leading PCI RAID controller at the RAID6 level of disk protection with a global “hot spare” disk. Since RAID6 keeps stripe parity on two disks, each ExaGrid appliance can tolerate the loss of up to two disk drives at the same time. The first lost disk drive will initiate a parity rebuild operation using the global hot spare as well as informing the backup administrator and (optionally) ExaGrid customer support of the failure. A replacement disk drive is dispatched quickly, typically allowing replacement of the failed disk the next business day. Loss of second disk does not result in loss of data since the remaining parity disk allows for data regeneration; this extends even longer the time available to replace the failed disk(s).

During normal operation, the RAID controller does consistency checking of the data on its disks in the background, correcting any disk media errors using the parity disks.

Flash-Backed RAID Cache

The industry-leading PCI RAID controller has onboard writeback cache backed up by a super-cap powered flash memory. Unexpected loss of appliance power does not result in backup data loss because any in-process writes to any disk are preserved until power is restored.

Backup Data Checksums with Automatic Repair

As backup data is deduplicated, checksums are added to the deduplicated data as it is placed into the internal storage area, called the “repository.” These end-to-end checksums cover the deduplicated backup data itself, and are used to verify the backup data during processing and as it is read from disk. The deduplicated backup data can optionally be replicated to a remote site; these checksums are used to validate the replicated data as well.

The ExaGrid software continually scrubs the repository data, confirming checksums and automatically repairing any deduplicated data that does not match its checksum using data from remote site(s). This automatic repair of deduplicated data is covered by one of ExaGrid’s patents.

Deduplicated Metadata Transactional Consistency

Metadata that tracks all of the deduplicated data is kept in a database and on internal storage. Software techniques are used to ensure transactional integrity of all metadata changes, including flushing filesystem pages into the flash-backed RAID onboard cache. The data flow of deduplicated backup data is protected end-to-end by the combination of checksums (above) and metadata transactional consistency.

Internal Database Backups and Self-Describing Metadata

The database used to keep metadata that tracks deduplicated data is periodically dumped to internal storage. These dumps are used to quickly restore the metadata database in the case of massive failure. The database dumps are used as an optimization; the metadata kept on disk is self-describing and can be used to completely rebuild the deduplicated data in the internal repository both at the local and remote ExaGrid sites.

Logging Filesystem

Backup data is kept in the ExaGrid internal storage on an industry-standard logging filesystem where file activity is logged for integrity and quick repair after an unclean shutdown.

Data Security

The data security capabilities in the ExaGrid product line, including optional enterprise-class Self-Encrypting Drive (SED) technology, provide a high level of security for data at rest and can help reduce IT drive retirement costs in the data center. All data on the disk drive is encrypted automatically without any action required by users. Encryption and authentication keys are never accessible to outside systems where they can be stolen. Unlike software-based encryption methods, SEDs typically have a better throughput rate, particularly during extensive read operations.

Data can be encrypted during replication between ExaGrid systems. Encryption occurs on the sending ExaGrid system, is encrypted as it traverses the WAN, and is decrypted at the target ExaGrid system. This eliminates the need for a VPN to perform encryption across the WAN.

Active Directory Support

ExaGrid integrates with Windows Active Directory for centralized credentials management and authentication that can be used to authenticate and authorize access to the ExaGrid management interface and backup target shares.

Periodic Assessments Using a Network Vulnerability Scanner

A complete vulnerability assessment is run periodically against ExaGrid's software. Vulnerabilities flagged by this assessment are evaluated and tracked and mitigated as appropriate.

Offsite Data Protection for Disaster Recovery

ExaGrid appliances can easily maintain offsite backups through the use of an offsite ExaGrid appliance in conjunction with a primary site ExaGrid appliance.

Backing up your data to an ExaGrid appliance at your primary site dramatically reduces the amount of disk space required to store all of that data due to its high-performance data deduplication capability. In a multi-site ExaGrid environment, the onsite ExaGrid system is only sending deduplicated data—the backup data bytes that change between each backup—over the wide area network (WAN) to the offsite ExaGrid appliance. The offsite ExaGrid appliance is ready for data restore and fast recovery in the event of a disaster or other primary site outage.

If the replication is one way only, the second site / offsite ExaGrid can be half the capacity of the primary site ExaGrid greatly reducing overall cost.

Replication between ExaGrid systems across a WAN can be scheduled for the day of the week and multiple times throughout each day. Each scheduled period allows for bandwidth throttling which limits replication to only use the assigned bandwidth. The combination of scheduling flexibility and bandwidth throttling allows for the maximum efficiency of WAN bandwidth used for replication. Replicated data can be encrypted over the WAN using a customer's VPN or by utilizing the ExaGrid built-in replication encryption.

ExaGrid supports various DR options:

Private Cloud

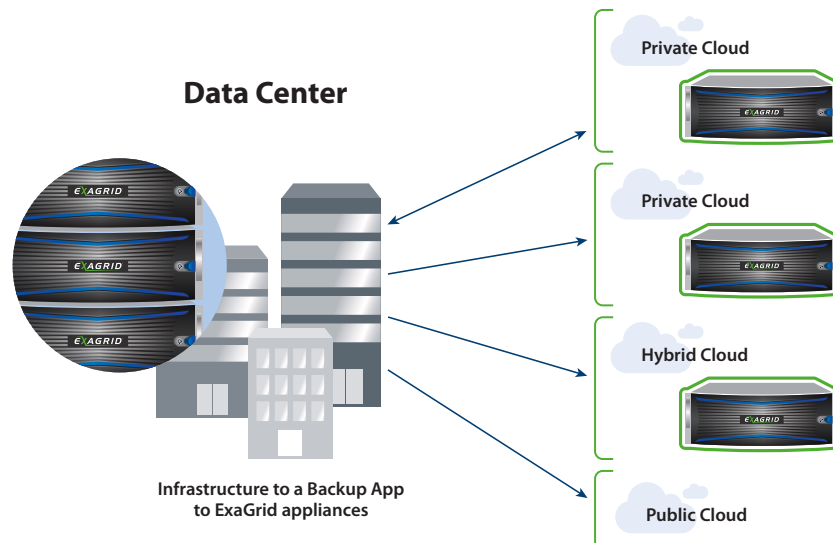
- Replicating to an ExaGrid at a customer's second data center (DR site)
- Replicating to an ExaGrid at a third-party hosted data center (DR site)

Hybrid Cloud

- Replicating to an ExaGrid owned and operated by an ExaGrid third-party DR provider or ExaGrid reseller and billed by the GB per month using OPEX budget

Public Cloud

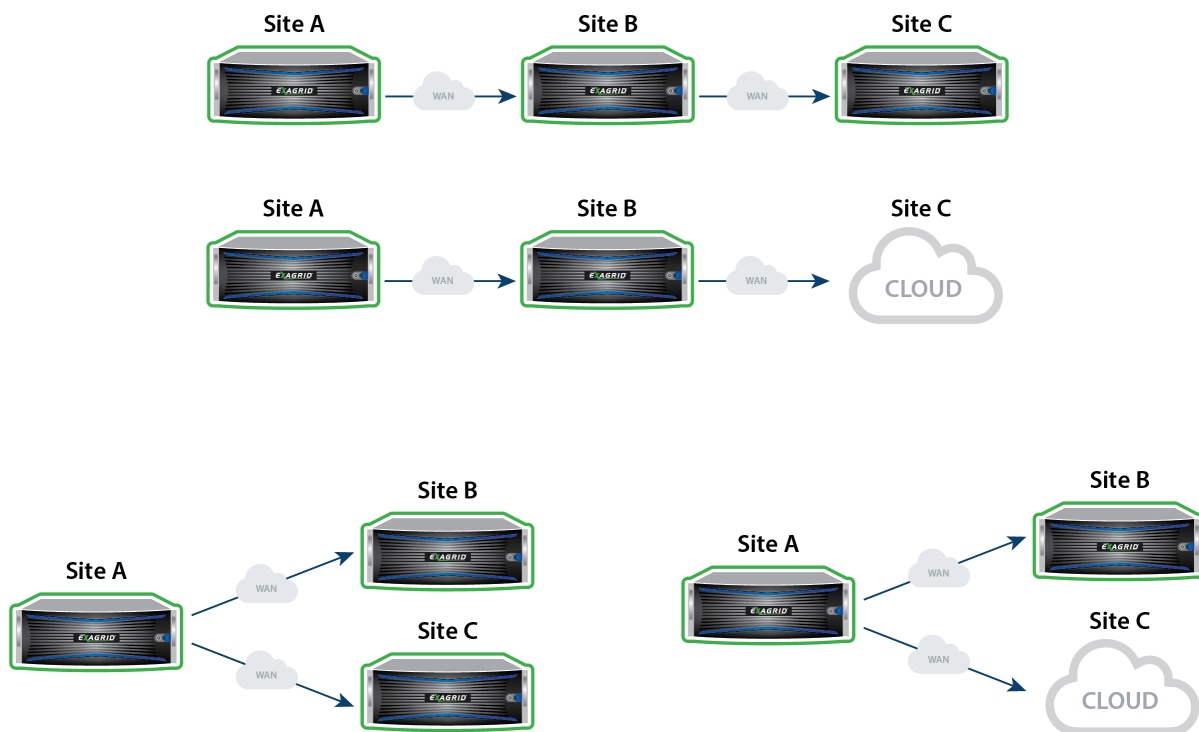
- Replicating to an ExaGrid VM in a public cloud (e.g., AWS), where DR data is stored in the public cloud and billed by the GB per month using OPEX budget




ExaGrid supports three models for private cloud DR sites at a customer's offsite data center:

- **Unidirectional replication to offsite for disaster recovery** – In this use case, the entire offsite system can be configured for repository, allowing for a half-size system to be used offsite. ExaGrid is asymmetrical in this use case where all other solutions are symmetrical.
- **Cross protection** – In this use case, data is backed up at both the offsite and onsite systems and cross replicated such that each site becomes the disaster recovery site for the other.
- **Multi-hop** – ExaGrid allows for a tertiary copy with two different topologies.
 - Site A can replicate to site B and then site B can replicate to site C
 - Site A can replicate to site B and site A can also replicate to site C
 - Site C can be a physical site or a cloud provider such as Amazon AWS
- **Multiple data center sites** – ExaGrid can support up to 16 sites in a single star topology with 15 spokes to a hub. Full systems or individual shares can be cross replicated such that data center sites can serve as disaster recovery sites for each other.

Multi Hop



Multiple Data Center



Total Cost of Ownership

Backup and disaster recovery is one area of IT spending which—though critically necessary—is typically viewed purely as cost. While backup is an extremely important area, organizations look to achieve appropriate protection so data is not lost while, at the same time, accomplishing this at the lowest possible cost. ExaGrid is the only vendor that has responded by creating a product that truly meets this different economic model warranted for backup spending. Backup spending has followed the same spending pattern as other IT infrastructure spending, which includes:

- Expensive forklift upgrades when a system is outgrown due to data growth
- Repurchasing of entire systems when an existing system simply “wears out”
- Complete rip-and-replace when a backup system becomes obsolete due to product end of life

ExaGrid redefines the economics of backup by helping you contain costs at every point in the life cycle — up front and as data grows over time.

ExaGrid offers the following to control costs:

- Over the phone/web installation at no charge – typically installed in a few hours
- No forklift upgrades
- No planned product obsolescence
- Scale-out architecture, pay as you grow
- For one-way replication – second site / DR site can be set to all repository, resulting in a system half the size at the DR site
- Everything is included in the yearly support and maintenance – no hidden cost
 - Local customer support in the Americas, EMEA and APAC – in theatre support
 - Assigned level 2 tech support engineer – work with the same person all the time
 - Automatic health monitoring
 - Full and point versions – all new features are included
 - No costs for failed hardware replacement
 - Spares depots around the world for fast failed hardware replacement
- Price protection
 - the price paid for appliances stays at the price for 5 years
 - yearly maintenance and support will not increase more than 3% per year

Cost Effectiveness Up Front

With ExaGrid Tiered Backup Storage, backups are written directly to a disk-cache Landing Zone, avoiding inline processing and ensuring the highest possible backup performance resulting in the shortest backup window. Adaptive Deduplication performs deduplication and replication in parallel with backups while providing full system resources to the backups for the strongest offsite recovery point (RPO). Available system cycles are utilized to perform deduplication and offsite replication for an optimal recovery point at the disaster recovery site. Once complete, the onsite data is protected and immediately available in its full unduplicated form for fast restores, VM Instant Boot and Recoveries, and tape copies while the offsite data is ready for disaster recovery. This allows ExaGrid to be more cost effective up front than an inline/block deduplication system.

In contrast, other appliances that use inline, block-level deduplication rely on a top-tier processor with large amounts of very fast memory and disk just to keep up with backup data. The premium cost of these components means higher cost compared to ExaGrid appliances. In addition, since the inline scale-up appliances have a fixed and limited ingest, in order to increase ingest, they need to use compute from elsewhere in the environment. To this end, they deploy software on media servers and certain application servers in order to do some of the deduplication work away from the inline appliance. Not only is deploying agents frowned upon by most IT organizations, but even with this approach, the ExaGrid will still be three times the ingest performance, and with ExaGrid there are no agents to deploy or manage.

When comparing ExaGrid appliances deduplication in the backup application software, it is important to keep in mind that using deduplication in the backup application software typically requires greater resources on the backup server—more processing power, more memory, and more disk. Software deduplication merely shifts the backup performance bottleneck to the media server. Using data deduplication in the backup software uses more disk and bandwidth over time and does not allow for backup environment flexibility such as using a separate utility for virtualized backup, direct TAR backups, and direct database dumps such as SQL dumps or Oracle RMAN dumps. ExaGrid's deduplication will be three to ten times more efficient. In addition, ExaGrid allows Veeam and Commvault deduplication to be turned on and ExaGrid with further deduplicate that data greatly increasing the deduplication ratio to save on storage costs.

Cost Effectiveness as Data Grows

ExaGrid's integrated Landing Zone with unique architecture—full appliances in a scale-out architecture is the most cost-effective way to scale as data grows. Each ExaGrid appliance added to the system includes a full server with additional processor, memory, bandwidth, and disk resources. Total backup capacity keeps pace with continued data growth over time by simply adding ExaGrid appliances to the system. There are no forklift upgrades and no additional future costs to consider.

Other appliances that use inline, block-level deduplication do not support a scale-out architecture and are therefore more costly to scale. Instead of adding capacity by adding full servers, only disk shelves are added over time as data grows. But, at some point, the single front-end controller becomes a bottleneck due to its fixed processor, memory and bandwidth resources and can no longer handle the backup load. Eventually, the entire front-end server must be replaced with the next higher capacity unit in a "forklift upgrade." In fact, you may have to spend as much for the front-end controller upgrade as you originally spent on the original system, including disk shelves. In addition, all data is always deduplicated. For each restore, recovery, and copy request, the data has to be put back together, or "rehydrated," which can take hours to days.

In addition, unlike other appliances that "end-of-life" in as little as 18 months and are incompatible with newer models from the same vendor, ExaGrid's scale-out architecture allows you to "mix and match" different capacities and generations of appliances within a single system. Only ExaGrid protects your backup investment from obsolescence.

Summary

When organizations evaluate the backup solution and company that can best meet their backup needs and address their challenges, more and more IT organizations are finding that ExaGrid offers not only the fastest backup and restore performance and best scalability but also a total cost of ownership that is typically at least half that of other solutions.

Only ExaGrid's unique scale-out architecture and Adaptive Deduplication provide:

Landing Zone Tier

- Fastest backups – avoids inline deduplication bottlenecks
- Fastest restores – no deduplicated data rehydration required

Retention Tier

- Low-cost long-term deduplicated retention storage
- Industry-leading 20:1 data deduplication
 - Global Deduplication
- Adaptive Deduplication
- Deduplicates and replicates during the backup window
- Strong offsite RTO and RPO
- Retention Time-Lock for Ransomware Recovery
- Non-network-facing tier
- Delayed deletes
- Immutable deduplication objects

Scale-out Architecture

- Scales to a 2.69PB full back up in a single system at up to 488TB/hr.
- Fixed-length backup window as backup data grows
- Eliminates fork lift upgrades of scale-up architectures
- Mix and match appliances – any age and any size
- No product obsolescence (no end of life of maintenance and support)
- 7 different capacity sized appliance models
- Scales as your data grows

Disaster Recovery Site options

- Can replicate offsite for DR
- DR site capacity is half the capacity of primary – asymmetrical
- Fewer appliances
- Lower cost
- Cross replication from site A to B and B to A
- Up to 16 sites in a hub and spoke topology
- Multi-hop – site A to B to C, site A to B, and site A to C
- Public Cloud DR site – Amazon AWS

Easy to Install Appliance Model

- Remote installation in a few hours

Redundancy

- RAID6 disk storage protection with a hot spare
- Hot swappable storage drives
- Can survive two simultaneous drive failures
- Redundant power supplies
- System runs if either power supply fails

Integrated Systems Management and Security

- Single user interface for all appliances in a system and across sites
- Active Directory for management interface and backup target security
- SNMP interface for integration with enterprise management apps
- Role-based access control
- Retention Time-Lock for Ransomware Recovery
 - Non-network-facing tier
 - Delayed deletes
 - Immutable deduplication objects
- Two-factor authentication
- Data encrypted at rest
- Data encryption while replicating over the WAN
- Security checklist makes it easy to apply best practices
- Data is checksummed to ensure data integrity
- System logging to external enterprise management applications
- Internal self-describing database

Backup Application Support

- Over 25 backup applications and utilities
- Supports heterogeneous backup application environments
- Veeam
 - SOBR – automation for job management
 - Data Mover
 - Improved backup performance
 - Improved security
 - Improved synthetic full performance
 - Veeam deduplication can be enabled and ExaGrid deduplicates further
- Veritas NetBackup
 - NetBackup Accelerator support
 - Reconstitute a full backup for Veritas NetBackup Accelerator
 - NetBackup OST support
 - NetBackup AIR support
- Commvault
 - Commvault deduplication can be enabled and ExaGrid deduplicates further
 - Spill and Fill support
- Oracle RMAN Channel backup support

Worldwide Distribution and Support

- Over 2,700 installed customers with tens of thousands of appliances
- Installed in over 40 countries
- Spares depots around the world
- Customer support around the world
- Included automatic health check system
- Assigned level-2 support engineer
- Work with the same senior level tech all the time

Programs

- Product price protection for 5 years
- Maintenance and support price protection – won't go up more than 3% per year

About ExaGrid

ExaGrid provides Tiered Backup Storage with a unique disk-cache Landing Zone, long-term retention repository, and scale-out architecture. ExaGrid's Landing Zone provides for the fastest backups, restores, and instant VM recoveries. The retention repository offers the lowest cost for long-term retention. ExaGrid's scale-out architecture includes full appliances and ensures a fixed-length backup window as data grows, eliminating expensive forklift upgrades and product obsolescence. ExaGrid offers the only two-tiered backup storage approach with a non-network-facing tier, delayed deletes, and immutable objects to recover from ransomware attacks. Visit us at exagrid.com or connect with us on LinkedIn. See what our customers have to say about their own ExaGrid experiences and why they now spend significantly less time on backup in our [customer success stories](#).



United States
United Kingdom
Singapore

100 Campus Drive / Marlborough, MA 01752 / (800) 868-6985
200 Brook Drive / Green Park, Reading, Berkshire RG2 6UB / +44 (0) 1189 497 051
1 Raffles Place, #20-61 / One Raffles Place Tower 2 / 048616 / +65 6808 5574

ExaGrid reserves the right to change specifications or other product information without notice. ExaGrid and the ExaGrid logo are trademarks of ExaGrid Systems, Inc. All other trademarks are the property of their respective holders.

©2022 ExaGrid Systems, Inc. All rights reserved.