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 undeduplicated 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.

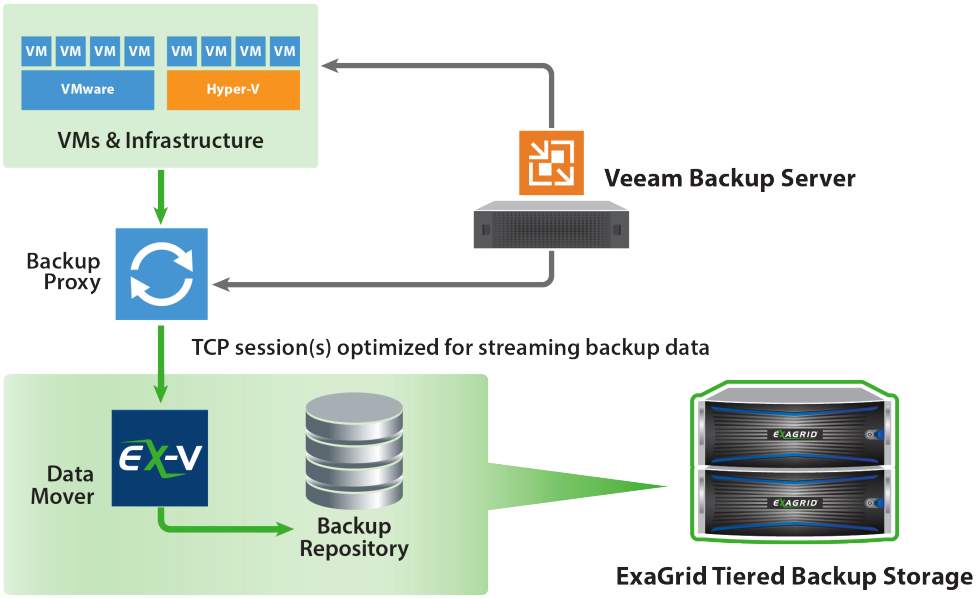

Veeam Scale-Out Backup Repository (SOBR)

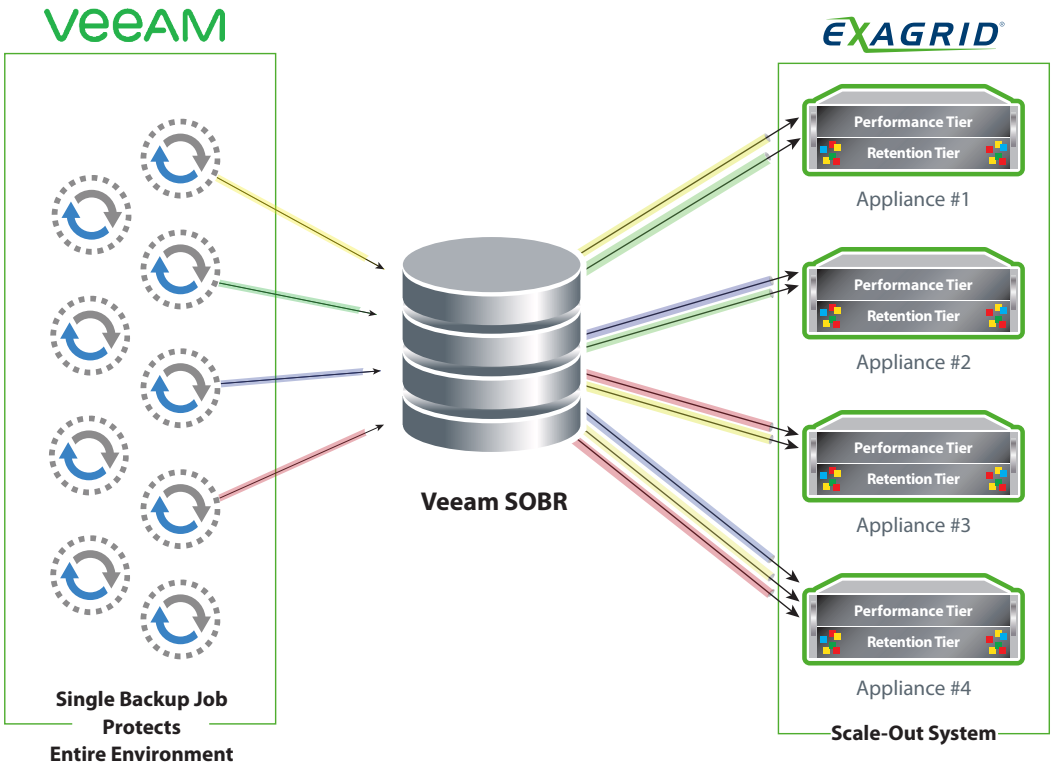

Veeam's Scale-Out Backup Repository (SOBR) allows backup administrators using Veeam to direct all jobs to a single repository made up of ExaGrid shares across multiple ExaGrid appliances with global deduplication in a scale-out system, automating job management to ExaGrid appliances. ExaGrid's support 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**

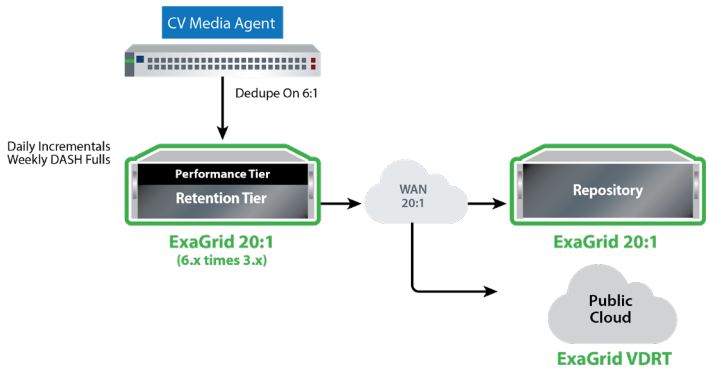

Commvault – **Dedupe Off**

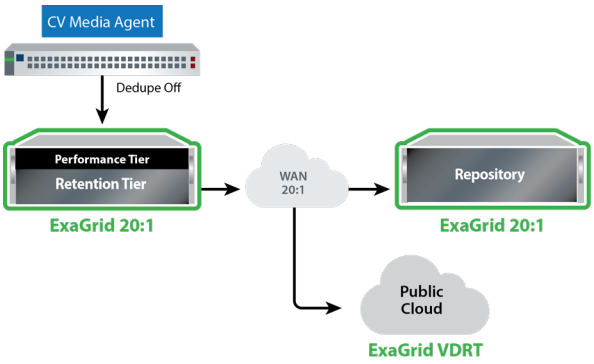

Commvault Dedupe On – **DASH Copy**

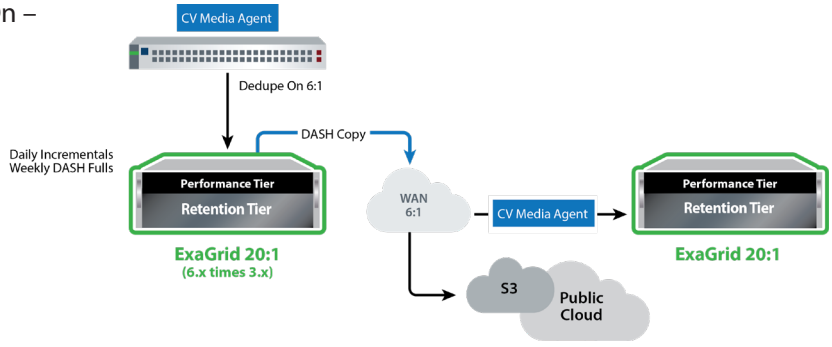

Oracle RMAN Channels

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

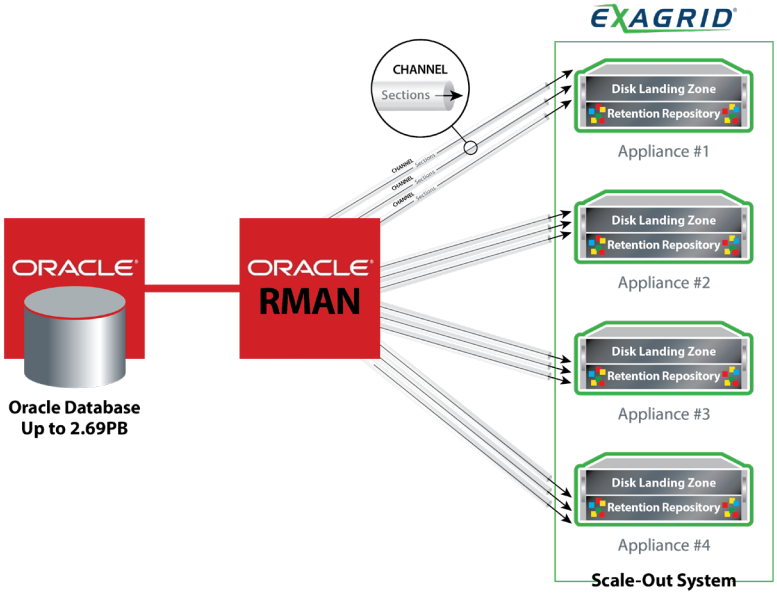

Reliability and Redundancy

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.

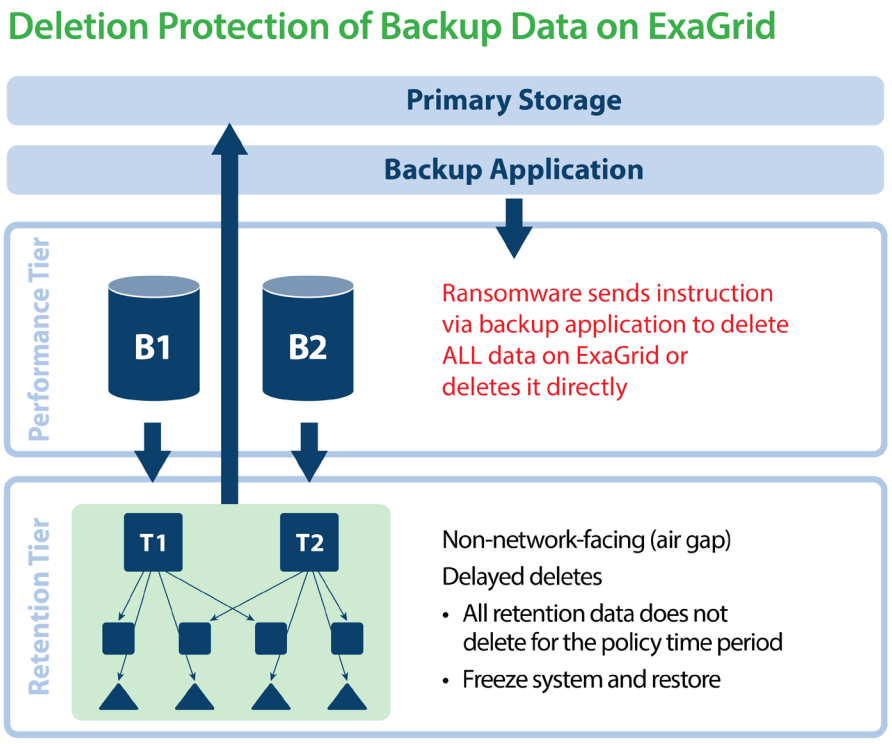

2) Data is encrypted in the ExaGrid disk-cache Landing Zone or is encrypted on the primary storage and backed up to ExaGrid such that ExaGrid has encrypted data in the Landing Zone and deduplicates it into the Retention Tier. The data in the Landing Zone is encrypted. However, all previously deduplicated 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

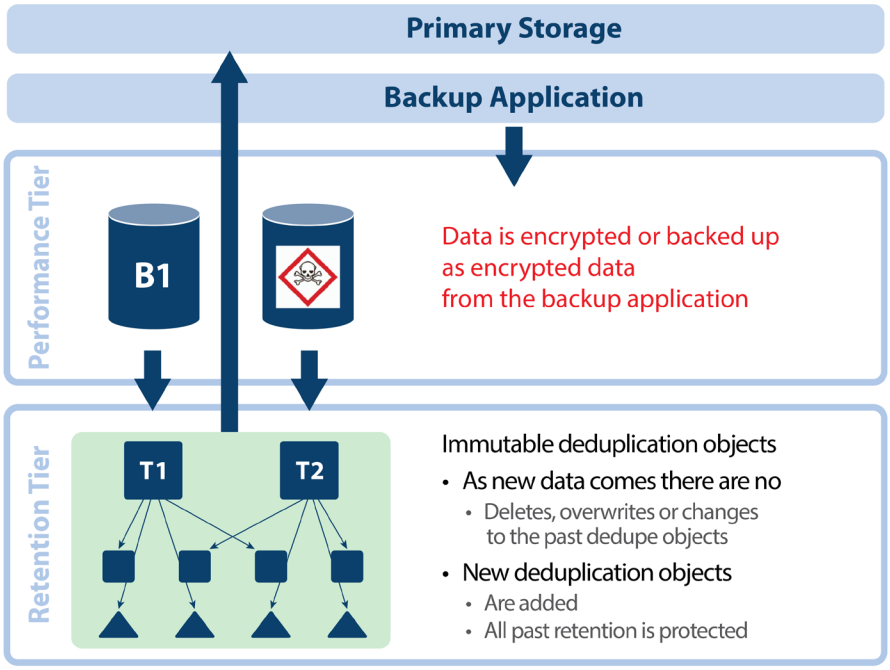

Features:

- 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

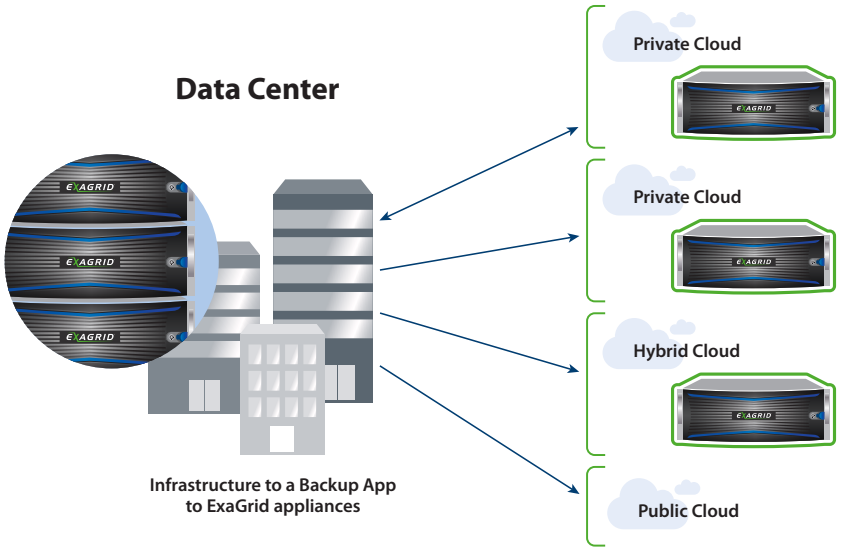

ExaGrid supports three models for private cloud DR sites at a customer's offsite data center:

- **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.

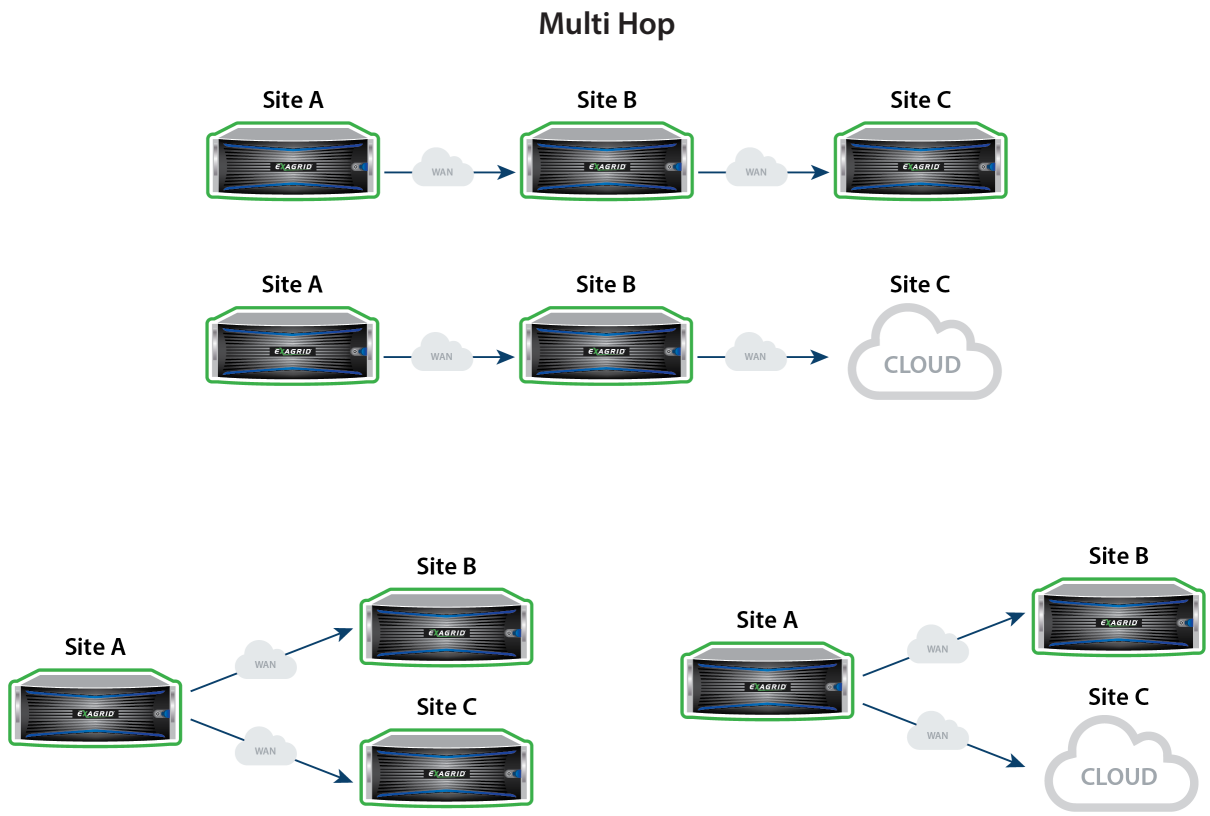

Multiple Data Center

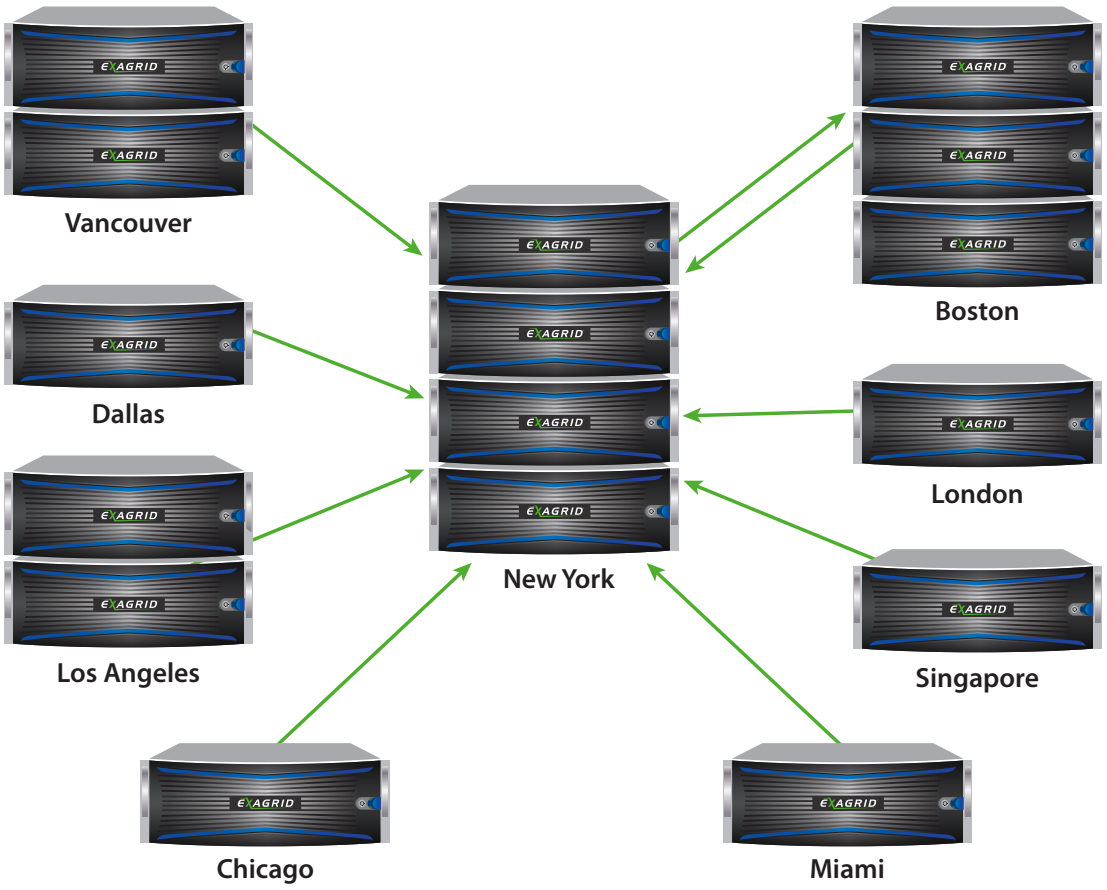

Total Cost of Ownership

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

★★★★★

ExaGrid Tiered Backup Storage

Fastest Backups.

Fastest Recoveries.

Unparalleled, Cost-effective Scale-out.

ExaGrid Product Overview

Tiered Backup Storage

ExaGrid's unique approach to backup storage delivers the fastest backups, restores, VM boots, and offsite tape copies as well as the only fixed-length backup window as data grows. In addition, ExaGrid's scale-out architecture and various size appliances allows customers to buy what they need as they need it, avoiding disruptive and costly forklift upgrades. Customers are able to mix older and newer appliances in the same scale-out system, eliminating product obsolescence and protecting their IT investment up front and over time.

Fastest Backups for the Shortest Backup Window

ExaGrid provides advanced and aggressive data deduplication, matching the high deduplication ratios in the industry of 10:1 to as high as 50:1 data reduction, with an average of 20:1, depending on retention periods and data types. However, ExaGrid understands that data deduplication is highly compute intensive and should not be performed during the backup window as the deduplication will slow down ingest performance and, as a result, will lengthen the backup window.

ExaGrid provides a unique disk-cache Landing Zone in each appliance where backups are written directly to disk so that the compute-intensive data deduplication process doesn't impact ingest speed. This approach provides the fastest backup ingest rate of any other deduplication solution. ExaGrid uses Adaptive Deduplication to deduplicate and replicate data to the disaster recovery (DR) site during the backup window (in parallel with the backups) but not inline between the backup application and the disk. This unique combination of a landing zone with adaptive deduplication provides for the fastest backup performance, resulting in the shortest backup window as well as a strong disaster recovery point (RPO).

Fastest Restores, VM Boots, and Offsite Tape Copies

Ninety-five percent or more of the total volume of restores, VM boots, and offsite tape copies come from the most recent backup, so keeping the most the most recent backup in only deduplicated form will require a compute-intensive, time-consuming data "rehydration" process that will slow down restore requests. VM boots can take hours from deduplicated data. Since ExaGrid writes directly to the disk-cache Landing Zone, the most recent backups are kept in their full, undeduplicated, native form. All restores, VM boots, and offsite tape copies are fast since the overhead of the data rehydration process is avoided. As an example, ExaGrid can provide the data for a VM boot in seconds to single-digit minutes versus hours for inline data deduplication backup storage appliances that only store deduplicated data. ExaGrid maintains all long-term retention (weeks, months, years) in a deduplicated format for storage efficiency.

Fixed-Length Backup Window

Since data deduplication uses a lot of processor and memory resources, as data grows, the amount of data deduplication to be performed grows as well. The first generation of deduplication storage appliances utilize a "scale-up" storage approach with a fixed resource front-end controller and disk shelves. As data grows, they only add storage capacity. Because the compute, processor, and memory are all fixed, as data continues to grow, so does the time it takes to deduplicate the data. The backup window becomes so long that the front-end controller has to be upgraded (called a "forklift" upgrade) to a larger/faster controller, which is disruptive and costly. Similarly, deduplication that is built into the backup software is far less aggressive, uses a larger amount of disk, and is much slower for backups and restores.

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.

Highest Performance for Backups

- Fastest backup performance for the shortest backup window by writing directly to a disk-cache Landing Zone, avoiding compute-intensive inline data deduplication.
- Backup windows kept permanently short as data grows by adding full servers (with processor, memory, disk, and bandwidth) in a single scale-out system.

Fastest Restores and VM Boots for Instant Recovery

- Fastest restore and tape copy performance from the most recent backup kept in its whole form. No reassembly from small blocks and large hash tables is required.
- Fast VM boots for instant recoveries from a high-speed Landing Zone, which maintains a non-deduplicated copy of the most recent backup. This approach avoids the time-consuming data rehydration required when using solutions that only store deduplicated data.

Most Cost-Effective Solution with No "Forklift" Upgrades

- Scalable next-generation architecture with full appliances provides plug-and-play expansion. To add an ExaGrid appliance, you simply plug it in and let ExaGrid's scale-out software virtualize the backup capacity pool.
- Multiple appliances allow full backups per appliance of 6TB, 10TB, 18TB, 27TB, 36TB, 52TB, and 84TB. Appliances can be mixed and matched with up to 32 appliances in a single scale-out system, allowing you to pay as you grow. Newer appliances can be added to older appliances in the same system to eliminate product obsolescence. With thirty-two 84TB appliances, a single system can support 5.37PB of usable storage and can ingest a 2.69PB full backup.
- 50% lower total system cost up front vs. competing systems. Over time, the total system cost is also 50% lower because the costly "forklift" upgrades associated with a first-generation front-end controller/disk shelf architecture are eliminated.

Advanced Features

- Scale-out architecture allows for cost-effective growth, eliminates product obsolescence, and maintains a fixed-length backup window as data grows.
- Unique Landing Zone reduces downtime by keeping a full copy of the most recent backup in complete form for instant recovery of VMs, full systems, and files. Competing solutions must rehydrate the most recent backup from millions or billions of deduplicated chunks causing much longer recovery time.
- Adaptive Deduplication performs deduplication and replication in parallel with backups while providing full system resources to the backups for the shortest backup window and an optimal recovery point (RPO) at the DR site.
- Plug and play expansion – various sized appliance models allow full backups of up to 84TB per appliance at an ingest rate of 488TB/ hour. Combining up to 32 appliances in a single scale-out system allows for scalability up to a 2.69PB full backup (5.37PB usable storage). In addition, ExaGrid supports second-site repository storage of up to 5.37PB for DR and long-term retention.
- ExaGrid includes replication to an offsite ExaGrid for disaster recovery, cross replication for multi-site DR and supports offsite tape copy creation.
- Private, hybrid, and public cloud DR support.
- Global deduplication across all appliances in a system.
- Bandwidth throttling for WAN efficiency.
- Management software notifies via SNMP or email that the system is reaching capacity thresholds.
- RAID6 guards against up to two simultaneous disk failures.
- Self-Encrypting Drive (SED) technology (encrypted models only) ensures that data at rest is always protected.
- WAN encryption for secure data transfer.
- Support of Oracle RMAN Channels for multi-hundred terabyte databases with automated performance load balancing and failover.
- Support of the Veeam Data Mover for synthetic fulls that are 6x faster.
- Support of Veeam SOBR for automated end-to-end scale-out backups to backup storage.
- Support of Veritas Backup Exec and NetBackup OST.
- Support of HYCU for Nutanix AHV and ESXi.
- A comprehensive list of over 25 supported backup apps and utilities can be found at **www.exagrid.com**.

United States 350 Campus Drive / Marlborough, MA 01752 / (800) 868-6985
United Kingdom 200 Brook Drive / Green Park, Reading, Berkshire RG2 6UB / +44 (0) 1189 497 051
Singapore 1 Raffles Place, #20-61 / One Raffles Place Tower 2 / 048616 / +65 6808 5574

exagrid.com

ExaGrid® Americas
Customer Support and Maintenance

DATA SHEET

"Best-in-Class" Disk Backup Solution in Under $100k and Under $50k 2013 Buyer's Guide Reports

ExaGrid Named 2013 "Disk Backup Champion"

InfoWorld.com Awards ExaGrid "Technology of the Year - 2013"

ExaGrid Wins "Disk Based Product of the Year: Small/Mid-range"

ExaGrid Recognized as "Top Emerging Vendor" in Customer Interest

ExaGrid's customer support and maintenance services are designed to ensure that ExaGrid meets your data protection needs.

Support and Maintenance

Annual Fee:	Yearly options based on a percentage of the actual purchase price of the system
Coverage:	All hardware, software and support coverage listed below

Support Response

Requirement:	Current annual customer maintenance and support renewal
Support Hours:	8:00 a.m. to 5:00 p.m. (Eastern standard time), Monday – Friday (optional 7x24 support available for an additional fee)
Methods:	Phone or email support
Response Time:	80% of phone calls and emails will be responded to in less than an hour

Installation

Service:	Installation is done via phone using a WebEx session. ExaGrid has installed thousands of customers and systems worldwide using this approach.

Self Monitoring, Automatic Notification, and Remote Support

Requirement:	Current annual customer maintenance and support renewal Valid remote access from an ExaGrid service center to the ExaGrid system
Service:	Monitor any alerts including pre-defined thresholds Remotely analyze and diagnose problems
Reconciliation:	Many problems are quickly resolved without customer intervention. ExaGrid does not commit to what percentage of problems it can resolve without customer intervention.

Hardware Maintenance

Requirement:	Current annual customer maintenance and support renewal
Program:	All systems are modular, and all drives and power supplies are hot swappable. Any failed hardware components are shipped next-day business air and are replaced by the customer. 100% of the hardware is covered—disk drive, power supply, server, included network components.

Software Maintenance

Requirement:	Current annual customer maintenance and support renewal
Program:	All versions (point and full) included at no charge. There are no additional charges.

Availability

Countries:	ExaGrid supports all countries in North, Central, and South America.

Contact Information

Support Email Address:	support@exagrid.com
Support Phone Number:	1.800.868.6985 or 1.508.898.2872 option 2

Using Veeam® Backup and Replication™ Software with an ExaGrid System

Veeam Version 10

ExaGrid Version 6.0 and higher

PN:3210-0500-02

Preface

This guide provides instructions on how to use the Veeam backup product with the ExaGrid System. This guide assumes the reader is experienced with Veeam and the ExaGrid System.

Veeam provides a number of ways to backup data. This guide provides basic backup examples. Experienced users will be able to use the information provided in this guide and build more complex backup schemes.

Purpose

The purpose of this guide is to suggest configurations, processes and procedures that will optimize the use of Veeam with the ExaGrid System.

Related Documentation

Title	Part Number
ExaGrid Administrator's Guide	210-0165-xx
ExaGrid Release Notes	210-0174-xx

Contacting ExaGrid Technical Support

If you have technical questions about this product that are not answered in this document contact ExaGrid Support:

- Email: support@exagrid.com
- Phone: 800-868-6985 and at the prompt, press 2.

Contents

Introduction

Note – This document applies to ExaGrid software version 5.2 and assumes the use of Veeam's Scale Out Backup Repository (SOBR) feature with the ExaGrid-Veeam Accelerated Data Mover.

The combination of ExaGrid's and Veeam's industry-leading solutions allows customers to take advantage of Veeam Backup & Replication on ExaGrid's disk-based backup system. This combination provides fast backups and efficient data storage as well as replication to an off site location for disaster recovery.

The ExaGrid system fully leverages Veeam Backup & Replication's built-in backup to disk capabilities and ExaGrid's zone-level data deduplication for additional data reduction (and cost reduction) over standard disk solutions.

Customers can use Veeam Backup & Replication's built-in source-side deduplication in concert with ExaGrid's disk-based backup system with zone-level deduplication to further shrink backups.

In addition, customers can replicate backups to off site storage for disaster recovery purposes without sacrificing performance of critical backup and recovery features.

The combination of Veeam with ExaGrid ensures fast, easy, and reliable backups, restores, and disaster recoveries.

ExaGrid/Veeam Overview

The ExaGrid System architecture is described in detail in the *ExaGrid Administrator's Guide* and ExaGrid online help. The following architectural information may be of particular interest to Veeam users:

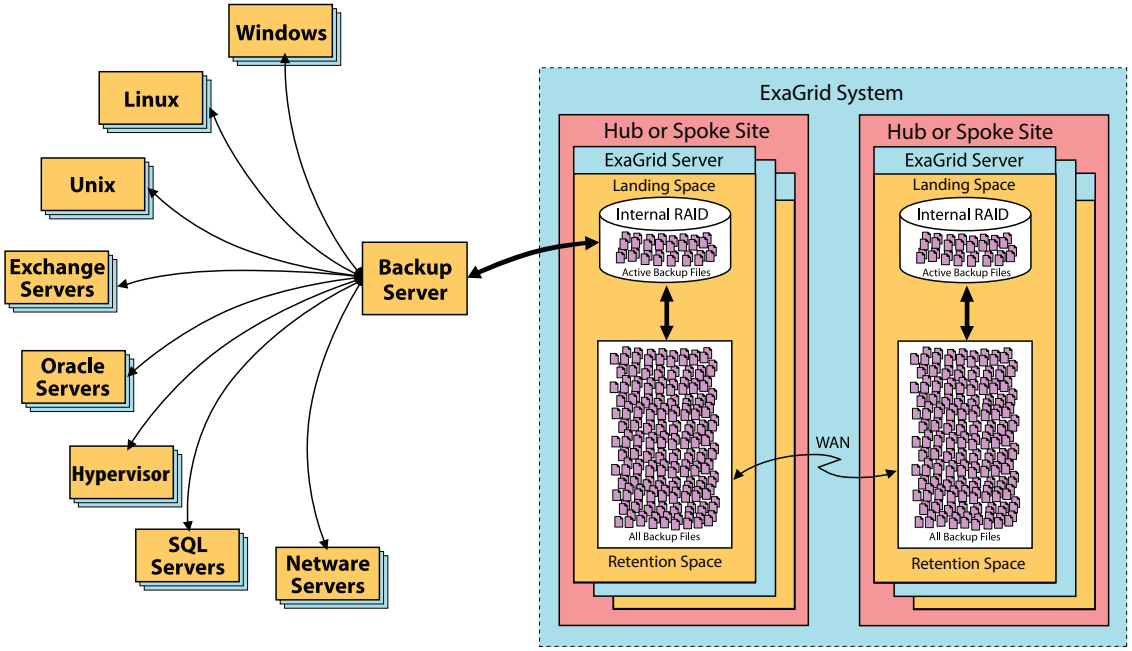

Data flows into an ExaGrid Share:

- Backup application writes to an ExaGrid share.
- When the backup is complete, the ExaGrid deduplication engine will deduplicate the backed up data.

Note – As more and more backups land, retained and deduplicated, deduplication ratios will improve.

- A full copy of the most recent backup is maintained in the Landing Space.
- The deduplicated versions are kept in the Retention Space.
- In multi-site ExaGrid Systems, you can specify whether or not to maintain a copy of a share on a second ExaGrid Site.

When a backup application requests a restore:

- Since most restores are done from the most recent backup, the ExaGrid system keeps the most recent backup in its entirety. This ensures extremely fast restores.
- Restores from an older backup are done by re-assembling the backup from backup versions.

ExaGrid's Landing Space and Veeam

ExaGrid's unique landing space architecture lends itself to Veeam's data protection features. The ExaGrid landing space is a high-speed disk cache that retains the most recent backups in complete form.

The following Veeam features take full advantage of ExaGrid's landing space.

Instant VM Recovery from an ExaGrid System

In the event of a primary storage outage or to prove your backups are recoverable, Veeam can instantly recover a virtual machine and run it directly from an ExaGrid Server's landing space. You then have the option to migrate the running VM to your production storage.

Instant File-Level and Application Item Recovery

Quickly recover an entire VM, an individual file from an image-level backup, or an Application Item Recovery. Veeam instant file-level recovery enables you to quickly restore individual files from your backups and replicas, taking advantage of ExaGrid's landing zone architecture to achieve the fastest restores and recoveries.

Veeam Explorers can be used to restore application specific items such as Exchange mailboxes, folders, Active Directory Users and Groups, databases and database objects.

Cloud Tier Copy Mode

Veeam can copy backups from the ExaGrid Landing Space to a cloud or object repository. For additional backup data protection, you can optionally set a retention lock on these copies.

See Veeam's documentation for more details about these Veeam features.

Veeam Long-Term Retention

Retention of backup for longer periods like weeklies, monthlies, yearlies etc. is simply configured in Veeam Backup Jobs via choosing a "GFS retention policy" See "Creating a VM Backup Job - Step-by-Step" on page 33.

Remote Backup Copy

Using Veeam's Backup Copy job feature to copy a backup from a local ExaGrid Server's landing space to a remote ExaGrid Server's landing space provides:

- A Veeam catalog that is "aware" of what is on the local and remote ExaGrid Sites
- Instant recovery operations from the landing space on either the local or remote ExaGrid Site
- Different retention definitions for the backups stored on the local and remote ExaGrid Sites

A Veeam Backup Copy job can use the ExaGrid-Veeam Accelerated Agent to copy the backups directly between ExaGrid sites, or a Veeam Backup Copy can take advantage of Veeam's WAN Accelerator to further reduce bandwidth required between the local and remote site.

Note – Do not confuse Veeam's Backup Copy to a remote ExaGrid Site with Veeam's VM replication feature or ExaGrid's replication feature.

For details on how to use Backup Copy with an ExaGrid System, see "Creating a Remote Backup Copy Job" on page 58

Tape Copy

With Veeam you can create backups that first write to an ExaGrid Share, then automatically copy the backup from the ExaGrid landing space off to tape.

This feature is especially useful to those ExaGrid customers who:

- Due to compliance and regulatory requirements, must maintain a removable copy of backed up data.
- Only have one ExaGrid Site and want to maintain a disaster recovery copy of their backups.

For details on how to use Veeam's Tape Copy with an ExaGrid System, see "Creating Backup Jobs for Physical Devices and Applications: Step-by-Step" on page 42.

ExaGrid-Veeam Accelerated Agent

The ExaGrid-Veeam Accelerated Agent is integrated with ExaGrid appliances and allows backups, restores and recoveries to complete faster. The Veeam backup server more efficiently inter-operates with its own Veeam data mover using optimized Veeam communications versus generic CIFS. In addition, the entire Veeam synthetic full operation occurs on the ExaGrid server eliminating the need to move data between the Veeam backup server and backup storage. This greatly reduces the time to complete a synthetic full.

The Veeam-specific protocol used between Veeam components (proxy servers, Backup and Replication Server, etc) and the ExaGrid-Veeam Accelerated Agent provides additional protection against ransomware attacks since the ExaGrid Veeam shares are not accessible through CIFS or NFS protocols.

Accelerated Synthetic Backups

An Accelerated Synthetic backup is the ability to use the ExaGrid-Veeam Accelerated Agent to write one full backup to an ExaGrid System, then going forward, using incremental backups, followed by a periodic synthetic full.

The entire synthetic full operation occurs on the ExaGrid Server, eliminating the need to move data between the Veeam backup server and backup storage which greatly reduces the time to complete a synthetic full.

Additionally, by only writing a single full backup to the ExaGrid System, then periodically synthesizing full backups on the ExaGrid System from incremental backups, you significantly reduce the demands placed on your virtual infrastructure by your backups.

Accelerated Synthetic full backups also reduce the burden on your Veeam servers by eliminating the need for moving full backup data and reducing the backup window of your incremental backups and adds the benefits of frequent, full restore points.

Accelerated Remote Backup Copy operations

A Veeam Backup Copy Job that runs from one ExaGrid Site (typically local) to another ExaGrid Site (typically remote), populates the destination ExaGrid share with backup copies that are retained using a GFS rotation controlled by Veeam settings in the Backup Copy Job. When the Veeam Backup Copy Job runs from a local ExaGrid to a remote ExaGrid, you achieve both Veeam's 3-2-1 rule (third copy)

and a retention of off-site backups that is different from retention of local backups - either shorter or typically longer. By using Veeam's WAN Accelerator, you can reduce the WAN bandwidth required.

Veeam's Scale Out Backup Repository

Veeam's Scale Out Backup Repository (SOBR) feature allows the storage capacity of all the ExaGrid Servers in a Site to be combined into a single Veeam backup storage repository that is used for all Veeam backups.

Using Veeam's Scale Out Backup Repository with ExaGrid has the following benefits:

1. Backup administrators no longer have to manually divide the Veeam backup jobs across multiple ExaGrid Servers in an ExaGrid site.
2. When Veeam starts a backup, Veeam selects the ExaGrid Server to which it sends backup data. Veeam makes this selection based on a policy setting that favors best ExaGrid deduplication and the free backup storage space available on the ExaGrid Site's ExaGrid Servers.
3. Moving a VM from one Veeam backup job to another Veeam backup job can be done without concern for which ExaGrid Servers are involved.
4. Fewer Veeam backup jobs need to be created/managed since a large Veeam backup job can simply target a Veeam SOBR repository and when the job runs, Veeam will select and target the optimal ExaGrid backup storage Server for each VM included in the job

SOBR significantly simplifies job configuration and management. Check with your Veeam representative on Veeam's Scale Out Backup Repository licensing considerations, as different Veeam versions have different limits on SOBR configuration.

Veeam's SOBR Performance & Capacity Tiers

ExaGrid Servers provide the Performance Tier of a Veeam SOBR configuration which can then be used as a source for Veeam operations to a Capacity Tier made up of cloud or other object storage. All Veeam Capacity Tier operations are supported with ExaGrid Servers used as the Performance Tier - such as SOBR offload of older backups or Cloud Tier Copy Mode.

Instant Recovery, Sure Backup and Application Item Recovery

These Veeam features can only be used from an ExaGrid Source Share. Do not use them with an ExaGrid InstantDR Share at a remote site.

Note – Though you cannot do an Instant Recovery from an InstantDR Share you can use Veeam to do a full VM restore from an InstantDR Share. To perform an Instant Recovery from a remote site, see "ExaGrid Veeam Share Migration" on page 70.

Guest File Level & Application Item Recovery

Can be used from either an ExaGrid Source Share or an ExaGrid InstantDR Share.

Importing Backups

Veeam's Import function allows access to Veeam backups on an InstantDR Share. Once imported, you can browse to and restore VMs and individual files from an ExaGrid InstantDR Share.

Upgrading from Veeam 9.5

Long-term retention of Veeam backups was typically done in Veeam 9.5 using a Backup Copy Job to the same ExaGrid Shares used for Veeam Backups.

The Backup Job produced the shorter retention (number of restore points) and the Copy Job produced the longer-term retention using a GFS scheme of weekly, monthly, yearly, etc.

Since long-term retention is now managed in Veeam V10 by a Backup Job setting, existing Veeam 9.5 customers should upgrade to Veeam V10 and migrate to Veeam V10's long-term retention feature:

1. Disable Veeam Backup Copy Jobs
2. Edit Veeam Jobs and enable the "**GFS retention policy**" and configure it to as needed.
3. Periodically use the Veeam console to manually delete the older backups that still exist on the repository used with the (disabled) Backup Copy jobs after they are no longer required.

   This step is needed because the Veeam Backup Copy Jobs are no longer running, and so the GFS rotation will no longer be automatically done by Veeam in the Backup Copy Job SOBR/Repositories.
4. After the oldest of these backups is deleted, the corresponding Veeam SOBR/extents and Backup Copy Jobs can be deleted from Veeam.

Managing Access

To use the ExaGrid-Veeam Accelerated Agent with an ExaGrid System, the ExaGrid Site to which Veeam will write must have an ExaGrid local user with credentials that match the Veeam server's credentials. See "ExaGrid-Veeam Accelerated Agent" on page 6 for more details about using the ExaGrid-Veeam Accelerated Agent.

Creating Local Users

To create an ExaGrid local user with credentials that can be used with Veeam:

Note – You must be logged into the ExaGrid System from a secure HTTPS connection otherwise the ExaGrid System will prompt you to re-log on.

1. From an ExaGrid System running version 6.0 or later, in the left-hand navigation frame click the ExaGrid Site that hosts the share to which Veeam will write.
2. In the top, main menu click **Site** and all the options available for the site are displayed.
3. Under **Security**, select **Local Users** and the Manage Local Users page is displayed.

4. Click **+New** and the Create Local User dialog box is displayed:

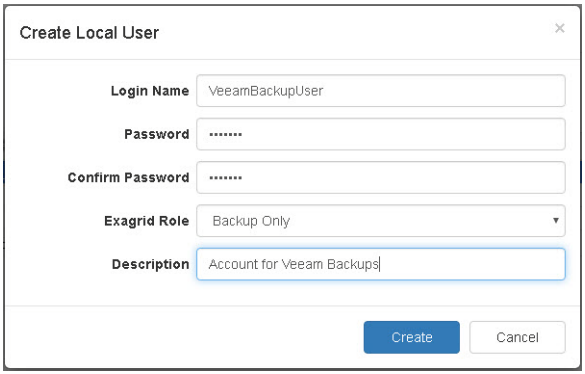

5. In the spaces provided, enter a user name and password.

Note – The user name and password are case-sensitive; they must exactly match credentials defined in Veeam.

---

6. From the **Management Role** drop down menu, select **Backup Only**.

Note – If the ExaGrid Site is included in an Active Directory Domain, the management role can be mapped into a Active Directory Domain. See the ExaGrid online help or the ExaGrid System Administrator's guide for details on using Active Directory Domains with ExaGrid Sites.

---

7. Click **Create** and the local user with the proper Veeam credentials is created.

Active Directory

Active Directory users and groups can be used to define access to an ExaGrid Share. See "User Access Policies" on page 12 for details.

The use of Active Directory with ExaGrid is described in detail in the ExaGrid Administrator's Guide and the ExaGrid Online help.

Caution – If you use an Active Directory user when adding a Veeam repository you MUST enter the user's Active Directory name in lower case regardless of how it was entered into Active Directory.

User Access Policies

User Access Policies define which users have access to a share.

The user's created with Veeam matching credentials must be added to an existing or new User Access Policy. The User Access Policy must then be assigned to the shares to be access by Veeam.

To define a User Access Policy:

1. In the left-hand navigation tree, select the ExaGrid Site on which you want to create the User Access Policy.
2. In the top, main menu click **Site** and all the options available for the site are displayed.
3. Under **Security**, select **User Access Policies** and the Share Access page is open to the User Access Policy tab.
4. Click **New** and the New User Access Policy dialog box is open.
5. In the space provided, enter a name for the User Access Policy.
6. In the space provided, enter an ExaGrid local user name, an Active Directory user name, or an Active Directory group name. See the ExaGrid Online Help or the ExaGrid Administrator's guide for details on using Active Directory.
7. Click the **+ plus** and the name or group is added to the Allowed User and Groups list.
8. Repeat the process for all users and groups that need access.
9. Click **Apply** and the User Access Policy is created.

ExaGrid Veeam Shares

Before you can use Veeam to backup to an ExaGrid System, you must first create an ExaGrid Share. You will then use Veeam to create and then write to a Veeam repository on the ExaGrid Share.

To take advantage of the ExaGrid-Veeam Accelerated Agent, when you create a new Veeam share, select the ExaGrid-Veeam Accelerated Agent protocol option.

The ExaGrid-Veeam Accelerated Agent option is particularly useful for those customers who want to take advantage of Veeam's:

- Synthetic full backups
- Backup Copy Job to a remote ExaGrid Site which enables Instant Recoveries and other Veeam vPower operations at the remote Site.

ExaGrid Share Considerations

Using Veeam's Scale Out Backup Repository greatly simplifies configuration and job management.

- If you are creating a share that will be used with a Veeam's Scale Out Backup Repository (SOBR) you only need to create one Veeam Share with ExaGrid-Veeam Accelerated Agent enabled on each ExaGrid Server in an ExaGrid Site. The Veeam SOBR will load balance backups automatically across all ExaGrid Severs, in the ExaGrid Site.
- Consider naming ExaGrid shares based on the ExaGrid Site name plus a unique suffix for each ExaGrid Server. For example, if the ExaGrid Site name is "Boston", and it has two ExaGrid Servers, name the ExaGrid Veeam share on the first server "Boston01", the share on the second server "Boston02", etc.
- Consider naming the Veeam SOBR that contains all the ExaGrid shares based on the ExaGrid Site name - eg "BostonBackupSOBR".

Creating a Veeam Share

This section provides a very high-level description of how to create an ExaGrid Share. Detailed instructions on how to create an ExaGrid Share can be found in the *ExaGrid Administrator's Guide* and the ExaGrid online help.

To create an ExaGrid Share:

1. Ensure that you have a local or domain user with proper Veeam credentials and that local user is added to a User Access Policy that can be used by this share. See "Managing Access" on page 10 for details.
2. In the left-hand navigation tree, click the ExaGrid Site's name on which you want to create the share.
3. In the top main menu, click **Site,** and all menu options for the site are displayed.
4. Under the Site's name, click **Shares and Replicas** and the ExaGrid Shares and Replicas page is displayed.
5. Click +**New**. The Create New Share dialog box is displayed.
6. In the space provided, enter a name for the share. Since there is only a single Veeam share required per ExaGrid Server when using Veeam's Scale Out Backup Repository, the naming convention must account for the number of ExaGrid

Servers in the ExaGrid Site. You must use unique share names on each ExaGrid Server. In addition, ExaGrid recommends using unique share names across your ExaGrid System. Doing so will:

- Make ExaGrid reports easier to interpret
- Avoid share name conflicts in disaster recovery scenarios
- Avoid share name conflicts if you decide to migrate a share to another ExaGrid Server.

7. From the drop down list, select the ExaGrid Server on which the share will be created.
8. From the **Type** drop down list select **Veeam backup & Replication**.
9. From the **Protocol** drop down list, select ExaGrid-Veeam Accelerated Agent. Select replication targets as needed.
10. From the **User Access Policy** drop down list, select the User Access Policy that contains the user created for this Veeam backup operation. If only one User Access Policy exists, that policy will be automatically selected.
11. Click **Create** and the Veeam Share is created. When using Veeam's Scale Out Backup Repository, only one share per ExaGrid Server is required for Veeam backups and Backup Copy Jobs. The share is combined in the Veeam Scale-out Backup Repositories with the other Veeam shares on the ExaGrid Site's servers.

Veeam Repositories

To write to an ExaGrid Veeam share, you must, from the Veeam use interface, create a repository. If you have two or more ExaGrid Servers, the multiple repositories can be combined into a single logical repository known as a Scale-out Backup Repository (SOBR).

Repositories use ExaGrid-Veeam Accelerated Agent credentials - see “Managing Access” on page 10. Ensure credentials have been created on the ExaGrid System prior to creating a Veeam repositories.

Repositories

Veeam repositories define the storage path for each Veeam Backup. A Veeam repository simply uses the root of the ExaGrid share.

Scale-out Backup Repositories

A Scale Out Backup Repository (SOBR) combines two or more repositories into a single logical repository that can automatically access the combined storage resources of the ExaGrid Servers that host the SOBR’s repositories.

Using a SOBR makes managing Veeam backup jobs easier because:

1. VMs can be easily moved from one Veeam job into another.
2. New VMs can be easily added to Veeam jobs. Consider using Veeam’s ability to include or exclude infrastructure folders in Veeam jobs as a mechanism to ensure newly create VMs are automatically included into the right Veeam backup job. See Veeam documentation for more details.
3. Jobs that have grown too large can be re-factored into multiple smaller jobs that continue to use the same ExaGrid SOBR backup repository.

Creating Backup Repositories

Using Veeam's Scale-out Backup Repository technology requires creating one share and repository for each ExaGrid Server in the ExaGrid Site. This will allow Veeam to take advantage of all of the resources of the ExaGrid Site.

Recommended Repository Settings

The following sections provide high-level descriptions of the required settings for creating Veeam/ExaGrid repositories. Experienced Veeam users can quickly review the settings and create repositories.

When creating a repository, for best results backing up to an ExaGrid Share:

- Considerations for the Veeam **Limit maximum concurrent tasks** repository setting:
  - Each ExaGrid server can support multiple concurrent Veeam tasks, and since the typical ExaGrid Site has multiple servers, the entire ExaGrid Site (accessed via single Veeam SOBR repository) can support a very large number of concurrent Veeam Tasks.
  - When creating a Veeam repository for an ExaGrid share, start with a **Limit maximum concurrent tasks** setting of 10 and work with ExaGrid Customer Support to further tune this setting.
  - For example, an ExaGrid Site with 8 ExaGrid Servers can initially support a total of 80 concurrent Veeam tasks and be tuned from there.
- Advanced Storage Compatibility Settings:

  Check **Use per-VM backup file**s. All other settings are optional and provide no deduplication advantage.
- Parallel Processing - Global setting for all repositories

  This option can be enabled for a Veeam server writing to an ExaGrid Share as long as the **Use per-VM backup files** option is set from the Advanced button on all repositories used with ExaGrid shares.

Creating a Repository - Step-by-Step

In this step-by-step example:

- Assume there are two (2) ExaGrid Servers in the ExaGrid Site.
- Create an ExaGrid Veeam Share on each ExaGrid Server.
- Create a Backup Repository on each of the two ExaGrid Shares.

To create a Backup Repository for an ExaGrid Veeam share.

1. From the Veeam left-hand navigation, click **BACKUP INFRASTRUCTURE**.
2. From the Veeam left-hand navigation, under BACKUP INFRASTRUCTURE, click **Backup Repositories**.
3. In the top-most menu, click **Add Repository** and the Add Backup Repository dialog box is displayed:

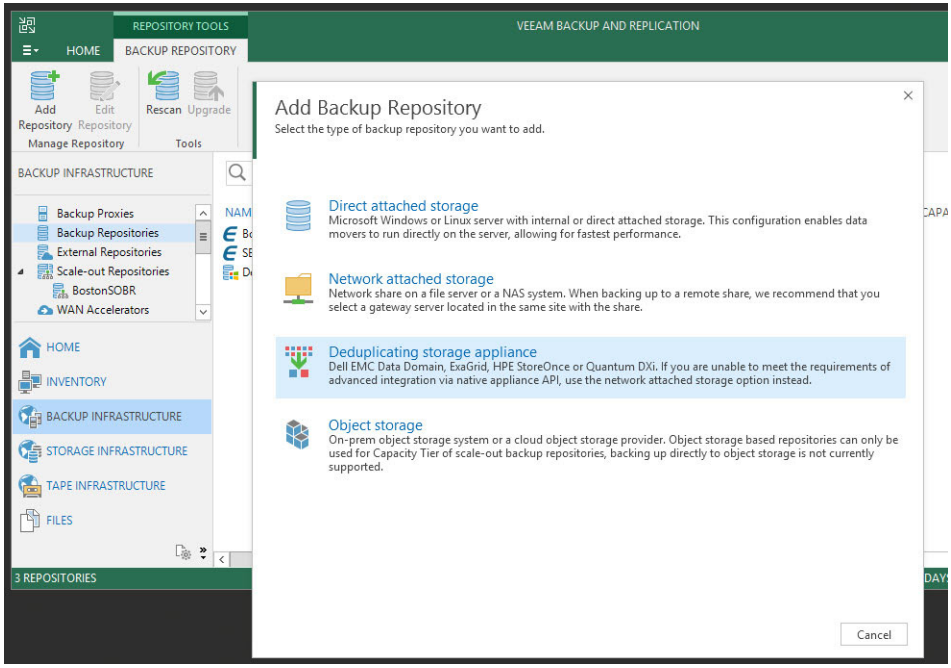

4. From the Add Backup Repository dialog box select Deduplication storage appliance and the Deduplicating storage appliances dialog box is displayed:

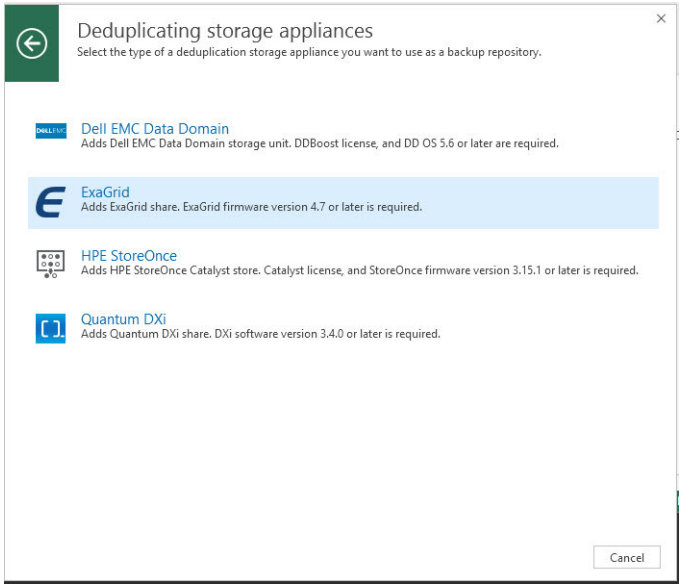

5. From the Deduplicating storage appliances dialog box select **ExaGrid** and the New Backup Repository wizard opens.:

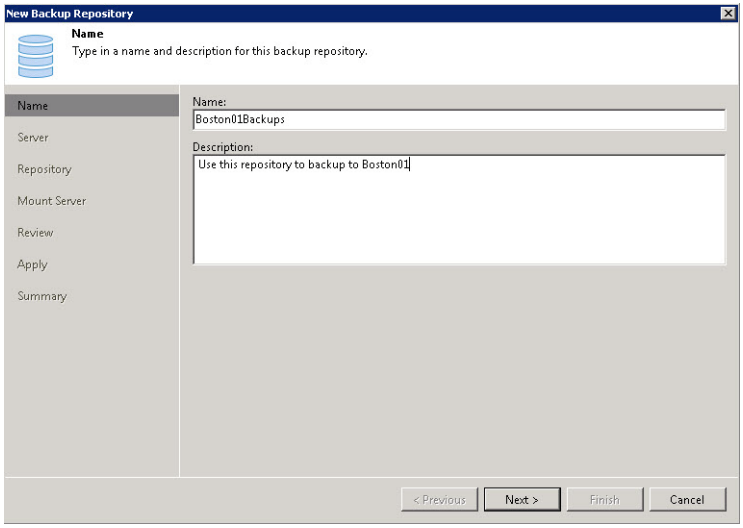

6. In the Spaces provided, enter a name for and a description of the backup repository.

7. Click **Next** and you will prompted for the Repository server information:

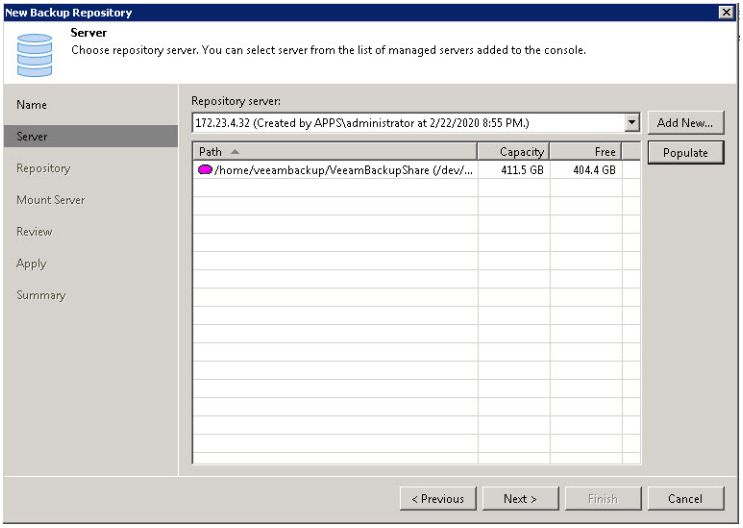

8. If you have not yet configured a Veeam Linux Server (which is typical), click **Add New...** and the Add New Linux Server dialog box is displayed.

   a. If you already defined a Veeam Linux Server (ExaGrid Server) for use with this repository, then from the Repository servers drop down menu select the server.

9. In the space provided enter the DNS name or an IP for one of the NICs for the ExaGrid Server that will become the Veeam Linux Server.

10. Enter a description as needed.

11. Click **Next** and the SSH Connection dialog box is displayed:

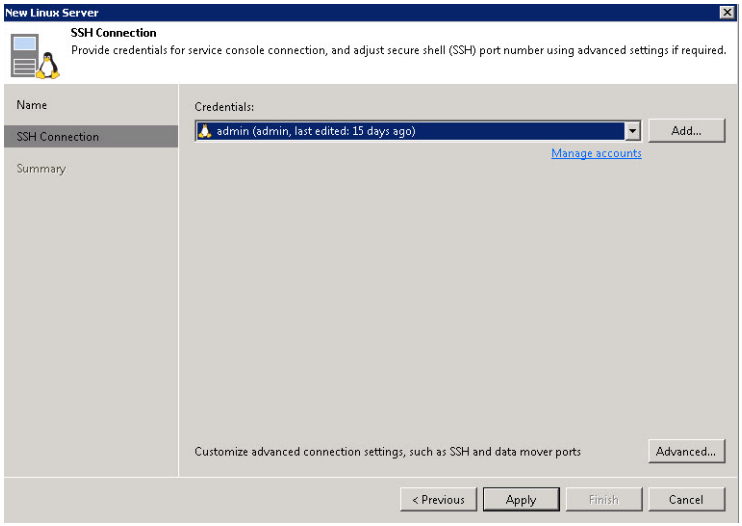

12. From the drop down menu can select a set of existing credentials or click **Add** to define a new set of credentials.

Note – Veeam Linux Server credentials are separate from SMB credentials. A Veeam server can require SMB credentials for CIFS shares as well as credentials for its managed Linux servers when using ExaGrid-Veeam Accelerated Agent shares.

---

13. If you are creating new credentials, in the spaces provided enter a user name and password pair that exactly matches the Veeam user credentials defined on the ExaGrid Server that will become the Veeam Linux Server.

Note – The credentials selected or defined for local Veeam users are case sensitive but Active Directory users must be entered all lower case.

---

14. Once you have selected or created the credentials, click **Advanced...** and the SSH Settings dialog box is displayed:

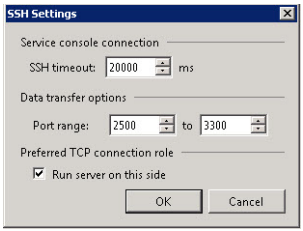

15. Check **Run this server on this side**. This ensures that Veeam will use the same ExaGrid network interface for restores and backups. This is especially important for those ExaGrid Systems with 10Gb or 40Gb network interfaces.

16. Click **OK** and the SSH Setting are set.

17. Click **Apply** and you will be prompted to review your settings.

18. Click **Finish** and the Linux server is created.

19. Once the Repository Server (ExaGrid Server) is selected or configured click **populate** and all of the ExaGrid Veeam Shares for that ExaGrid Server are listed

20. Click **Next** and the Repository dialog box is displayed:

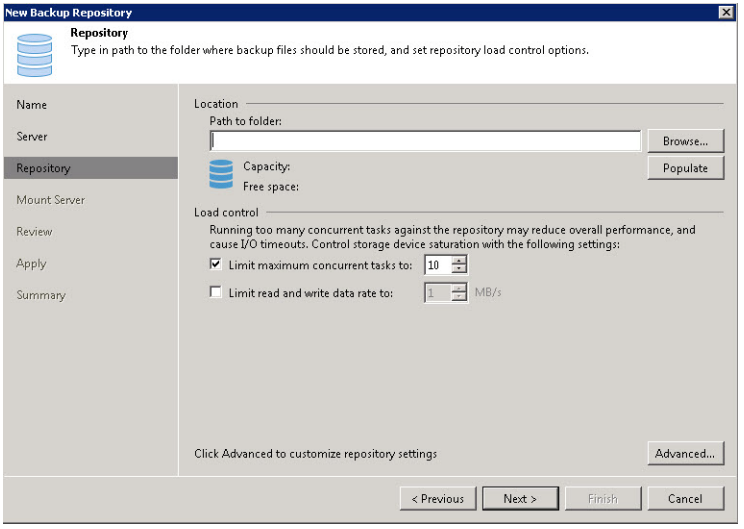

21. Click **Browse** and a list of folders is displayed

22. Expand the list of folders starting with the folder named “home/” and select the folder with the name of the ExaGrid Veeam Share you created for this repository.

Note – If a list of Veeam shares is not populated, confirm that the credentials used for this Veeam repository match exactly with those defined on the ExaGrid - see “Managing Access” on page 10

---

23. Select the Veeam Share you created.

24. Click **OK**. Confirm the “Path to folder” is correct:.

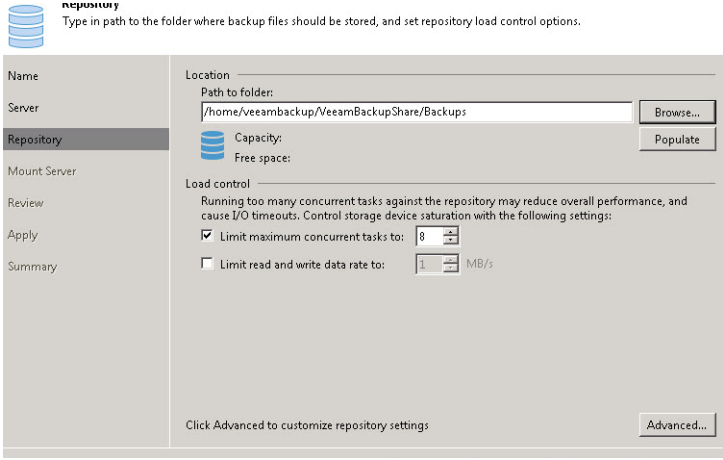

25. Set **Limit maximum concurrent tasks to:**
    - An initial value of 10 - see “Recommended Repository Settings” on page 17

26. Click **Advanced** and the Storage Compatibility Settings dialog box is displayed:

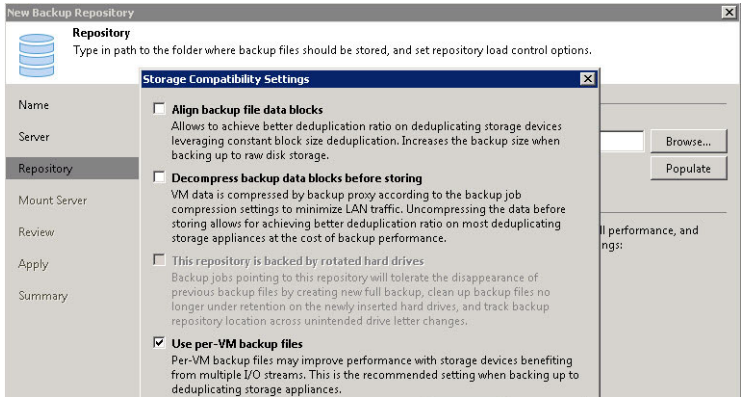

27. Ensure that **Align backup file data blocks** is unchecked. ExaGrid is a byte-level deduplication appliance and this option will not improve deduplication

28. Ensure that **Decompress backup data blocks before storing** is unchecked. Using this option will not significantly improve the additional deduplication achieved by the ExaGrid Server and the best practice is to leave this option unselected.

29. Ensure that **Use per-VM backup files** is checked. This maximizes the performance of backups and the deduplication achieved on the ExaGrid.

30. Click **OK** and the Repository information dialog box is re-displayed:

31. Click **Next**, confirm Veeam vPower settings, click **Next** again to make a final review of your settings.

32. Click **Finish** and the repository is created.

Once all of the repositories are created, create a Scale-out Backup Repository that includes all the repositories you created - thereby combining the backup storage capacity of the entire ExaGrid Site into a single Veeam SOBR.

Creating Scale-out Backup Repositories

Note – Some Veeam Editions limit the number of extents (ExaGrid Shares) that can be added to a SOBR repository. Consult with your Veeam representative on these limits.

Using Veeam's Scale Out Backup Repository makes Veeam job management significantly easier by letting Veeam decide where to put the per-VM backup files across multiple ExaGrid Servers.

Once a Veeam repository has been created for each ExaGrid Share on each ExaGrid Server, combine the repositories into a single SOBR repository. This enables Veeam to send backup data automatically to different ExaGrid Shares based on their available space. Veeam refers to each ExaGrid Share repository as an extent.

Before You Begin

Before you begin, you must create a Backup repository for each ExaGrid Server as described in "Creating Backup Repositories" on page 17.

You must create:

- One ExaGrid Share on each ExaGrid Server in the ExaGrid Site
- A repository for Backups on each ExaGrid Share

Recommended Settings

The following sections provide high-level descriptions of the required settings for creating Veeam/ExaGrid Scale Out Backup Repositories (SOBR). Experienced Veeam users can quickly review the settings and create a SOBR.

When creating a SOBR, for best results backing up to an ExaGrid Share:

- Some Veeam Editions limit the number of extents (ExaGrid shares) that can be added to a SOBR repository. Consult with your Veeam representative on these limits.
- You must create a Scale Out Backup Repositories for Backups Jobs. The SOBR will send Veeam backup data to the backup repositories on ExaGrid.
- Put all ExaGrid Servers that host shares to be used by the SOBR into the SOBR (each is a separate Veeam SOBR extent).

- In the **Advanced** settings, check both **Use per-VM backup files** and **Perform full backup when required extent is offline**.
- Use the Data locality policy

Creating a Scale-out Backup Repository - Step-by Step

In this step-by-step example:

- Assume the two (2) Backup repositories were created as described in “Creating a Repository - Step-by-Step” on page 18.
- Create a single Scale-out Backup Repository for the Backup repositories.

To Create a new Veeam Scale Out Backup Repository:

1. From the Veeam left-hand navigation, click **BACKUP INFRASTRUCTURE**.
2. From the Veeam left-hand navigation, under BACKUP INFRASTRUCTURE, click **Scale-out Repositories**.
3. In the top-most menu, click **Add Scale-out Repositories** and the New Scale-out Backup Repository dialog box is displayed:

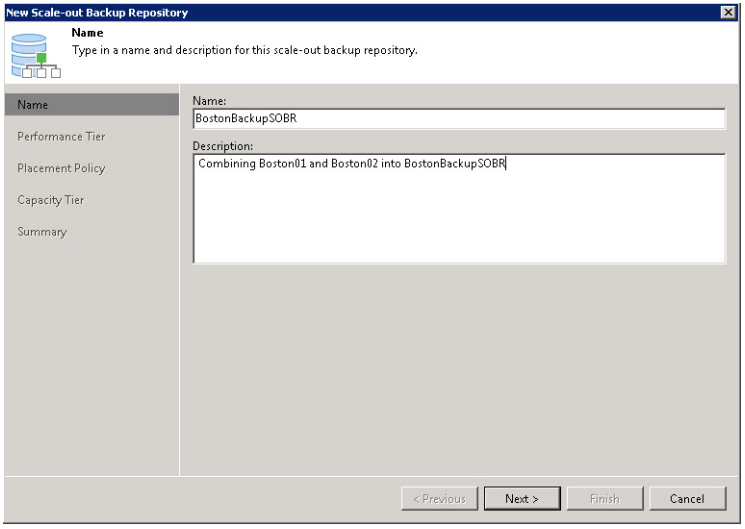

4. In the spaces provided enter a name and description for the SOBR.

   In this example we have named the SOBR so it’s name is easily identifiable as a SOBR made up of the backup repositories in Site Boston: BostonBackupSOBR.

5. Click **Next** and the Performance Tier options are displayed

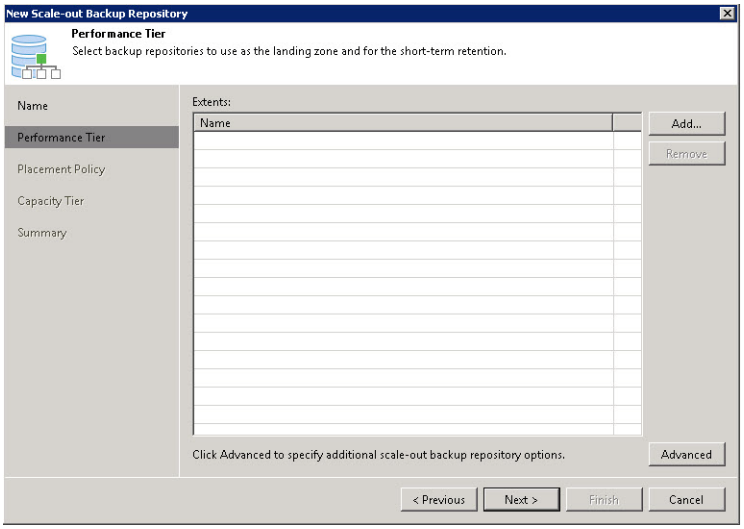

6. Click **Add** and the repositories (Extents) you created are displayed:

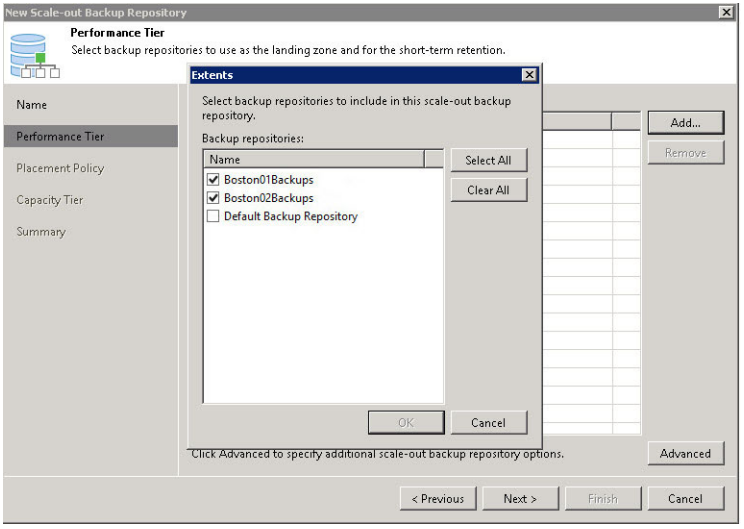

7. Select the repositories (Extents) that will be used by this SOBR.

8. Click **OK** and the repositories (Extents) will be listed in the Performance Tier:

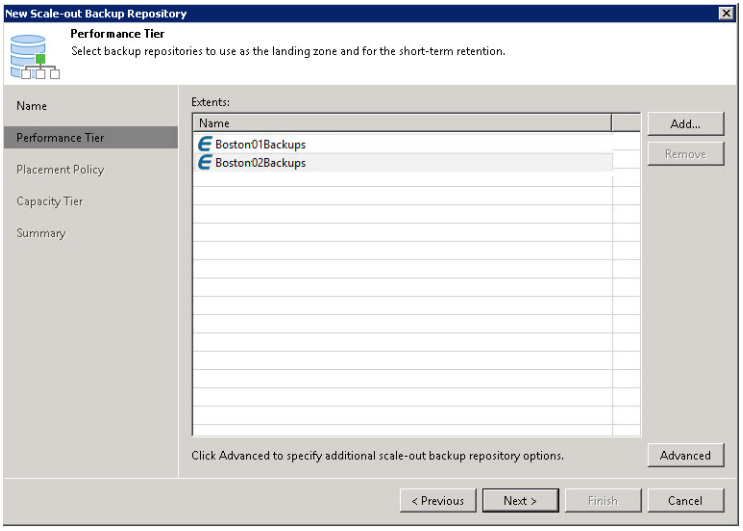

9. Click **Advanced** and the Advanced options for the SOBR are displayed:

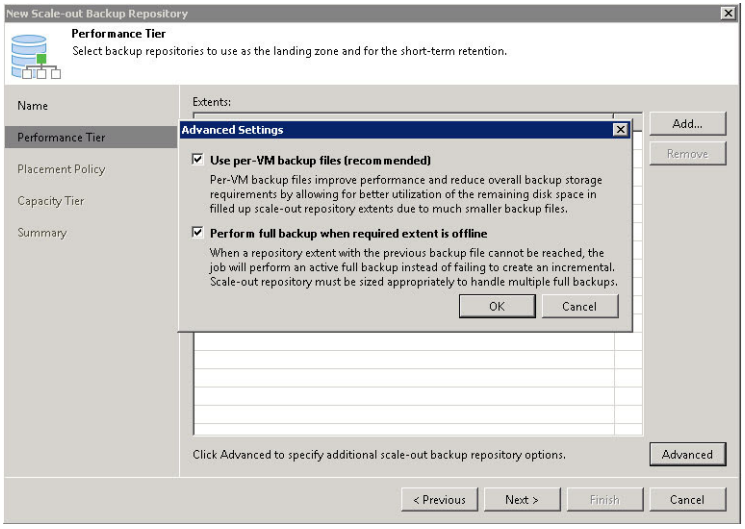

10. Check both **Use per VM backup files (recommended)**. and **Perform full backup when required extent is offline**. There may be cases where you would rather have a Veeam backup job fail due when a SOBR extent (ExaGrid Share) is offline - in which case disable the “Perform full backup when required extent is offline” setting.

11. Click **OK** and the Advanced options dialog box closes.

12. Click **Next** and the Placement Policy options are displayed:

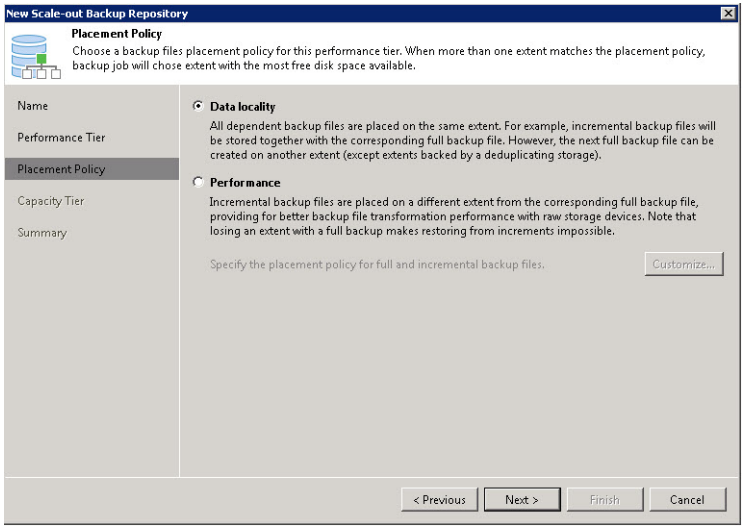

13. Select **Data Locality**.

14. Click **Next** and the Capacity Tier options are displayed.

15. Set Capacity Tier options as needed. When Veeam moves backups to the Capacity Tier they are replaced by stubs on the ExaGrid and will reduce ExaGrid Retention Space consumption. When Veeam copies backups to the backup Tier, they remain on the ExaGrid.

16. Click **Apply** and the SOBR is created and its status is displayed.

Veeam Backup Jobs - Overview

Before you can create a Veeam backup job you must

- Create at least one Veeam Share. For details, see “ExaGrid Veeam Shares” on page 13.
- Create at least one backup repository or SOBR repository. For details, see:
  - “Creating Backup Repositories” on page 17.
  - “Creating Scale-out Backup Repositories” on page 25.
- If you are backing up physical servers, create at least one Protection Group. See “Creating a Veeam Protection Group” on page 47.

Veeam allows you to backup both VMs and physical servers. VMs and physical servers use the same types of repositories and except where explicitly called out, use the same ExaGrid recommended settings.

See “Creating Backup Jobs for Physical Devices and Applications: Step-by-Step” on page 42 for a list of Veeam data sources supported by ExaGrid backup targets.

Recommended Job Settings

The following sections provide high-level descriptions of the required settings for creating Veeam/ExaGrid backup jobs. Experienced Veeam users can quickly review the settings and create jobs.

Backup Jobs using a SOBR

All jobs that send Veeam backup data to ExaGrid Servers need to only target the Veeam SOBR with all the Repositories (extents) created for backup on ExaGrid shares. As jobs run, Veeam will direct the per-VM backup files to separate ExaGrid Servers based on Veeam's policy setting and each server's free space.

Follow Veeam's recommendation on job limits and the proxies required to achieve multiple concurrent backup streams to the ExaGrid Servers.

Consider using Veeam's ability to include and exclude VMs by infrastructure folders as a mechanism to automatically include new VMs in Veeam backup jobs. Consult Veeam documentation for more details.

Backup Jobs

When defining retention:

- If more than 14 points or days are needed, consider using a Grandfather-Father-Son (GFS) policy.
- Consult with your ExaGrid customer support engineer if an extended retention plan is needed.

For the tabs in a Backup Jobs Advanced settings dialog box:

Backup tab:

- Ensure **Reversed Incrementals** is unchecked. Using this option negatively impacts ExaGrid deduplication and will likely cause consumption issue.
- Ensure **Incremental** is checked
- Ensure **Transform previous backups...** is unchecked.
- **Create synthetic fulls backups periodically**. See also "Synthetic Full Backup Jobs" on page 32 for details. Periodic synthetic fulls is a Veeam-recommended practice and will ensure the best ExaGrid deduplication and backup storage cost efficiency.
- If you do not enable synthetic fulls, schedule a periodic Veeam Active full job. Veeam's incremental forever model has a negative impact on deduplication when restore points limits are reached.
- When using synthetic fulls, schedule a monthly periodic active full as recommended by Veeam's best practices.

Maintenance tab:

- Ensure that **Perform backup files health check** is enabled for some regular interval. This periodic health check will not impact backup, restore, or deduplication performance.

- There is no need to do any **Full backup file maintenance**

Storage tab:

- Check **Enable inline data deduplication.**
- Set **Compression level** to **Dedupe-friendly.**

Caution – The two settings above are critical to achieving maximum overall deduplication and your ExaGrid solution has been sized assuming these settings. Deviating from this practice will likely result in ExaGrid capacity issues.

- Set **Storage optimization:** to **Local Target**.
- Ensure **Encryption** is not enabled
- Set other options as needed

For all other tabs, including the Integration tab, set options as needed.

Synthetic Full Backup Jobs

Synthetic full backups are written to the ExaGrid Server that hosts the Veeam Share. To minimize the impact on system performance, ExaGrid recommends running synthetic full backups according to Veeam's recommended recipe of: **After the initial full backup, run daily incremental and weekly synthetic full backups.**

Using the ExaGrid-Veeam Accelerated Agent will significantly reduce the time required to complete synthetic full backups.

Not all Veeam backup jobs support creating synthetic full backups - consult Veeam documentation for details.

Note – When using synthetic fulls, scheduled, monthly periodic active full backups are recommended.

Creating a VM Backup Job - Step-by-Step

Note – As a best practice, ExaGrid recommends using backups with a **GFS retention policy**. This allows Veeam to retain older backups longer while retaining the fewest number of Veeam Retention Policy Restore Points. For example, rather than keep a large number of daily backups, consider retaining 14 or less restore points (days of daily backups) and then with a GFS policy, retain multiple weeks, months, or years of backups.

To create a Veeam backup job:

1. From the **Home** tab in the Veeam user interface, click **Backup Job**.
2. From the drop down menu select **Virtual machine...** and the New Backup Job dialog box opens:

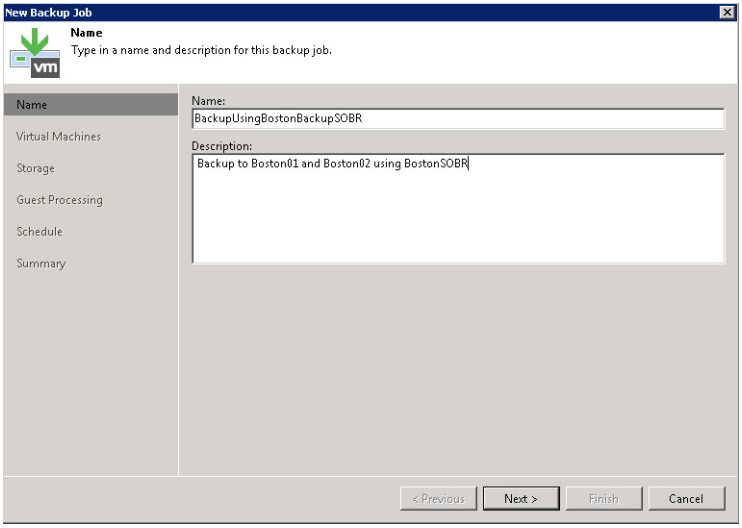

3. In the spaces provided, enter a name and description. Adapt a naming convention that easily identifies the backup job.
4. Click **Next** and the Virtual Machine options dialog box opens.
5. Click **Add** and the Add Objects dialog box opens.
6. Select the VMs to be backed up
7. Click **OK** and the Add Objects dialog box closes and the Virtual Machine option dialog will be populated with the VMs you selected:

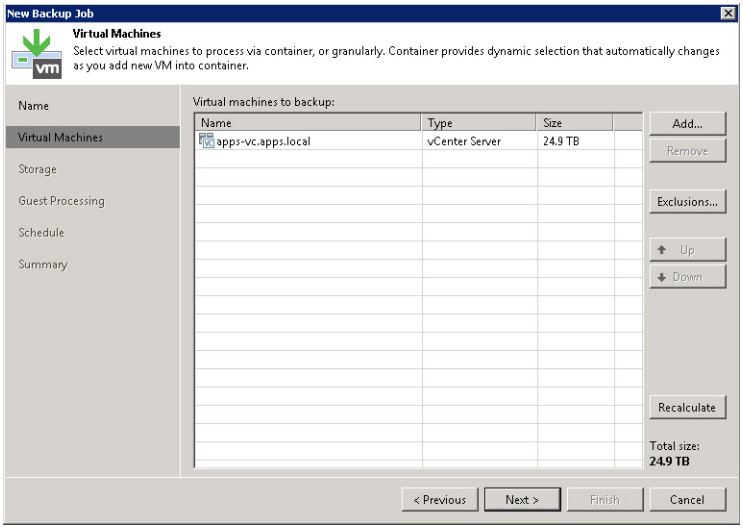

8. Click **Exclusions** to take advantage of Veeam's ability to include or exclude infrastructure folders. This feature allows newly created VMs to be automatically included into the right Veeam backup job. See Veeam documentation for more details.

9. Click **Next** and the Storage options dialog box opens:

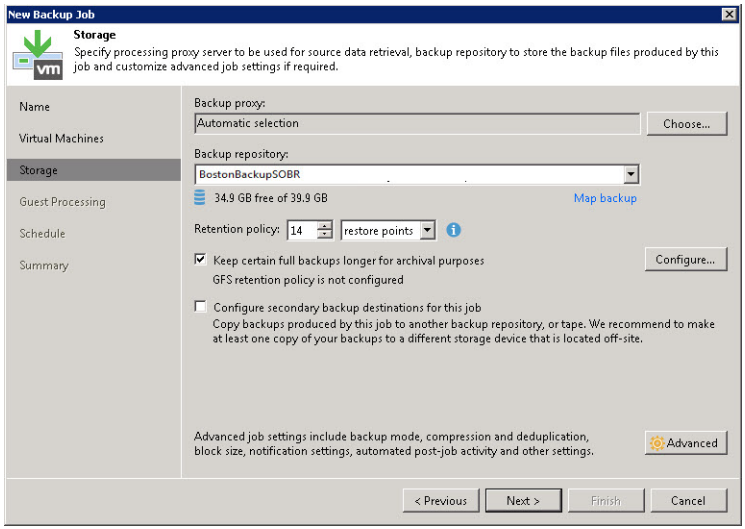

10. Next to the Backup proxy entry, click **Choose** and choose a backup proxy as needed.

11. From the Backup repository drop down list, select the repository you created for this backup job.

12. Set **Retention Policy Restore points to keep** as needed.

Note – As a best practice, ExaGrid recommends using a **GFS retention policy**. This allows Veeam to retain older backups longer while retaining the fewest number Veeam Retention Policy Restore Points. For example, rather than keep a large number of daily backups, consider retaining 14 or less restore points (days of daily backups) and then with a GFS policy, retain multiple weeks, months, or years of backups.

---

13. Configure a GFS (Grandfather-Father-Son) retention policy:
    - Check the **Keep certain full backups longer for archival purposes**.
    - Click **Configure...** and the Configure GFS dialog box opens:

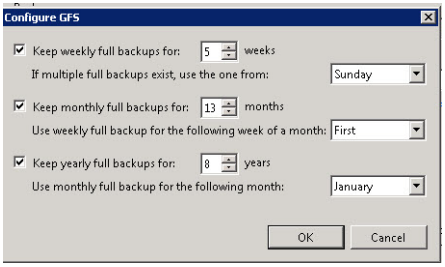

14. Configure the GFS Weekly, Monthly, and Yearly settings as needed. If your retention requirements have increased beyond the original ExaGrid sizing, please contact your ExaGrid Support Engineer and explain that your retention needs have changed.
    - Click **OK** and the Configure GFS dialog box closes.
15. Check **Configure secondary destination for this job,** if after the job completes, you want a copy of the backup written to tape directly from the ExaGrid landing space,
16. Click **Advanced** and you will be prompted to make advanced backup settings.

17. In the **Backup** tab:

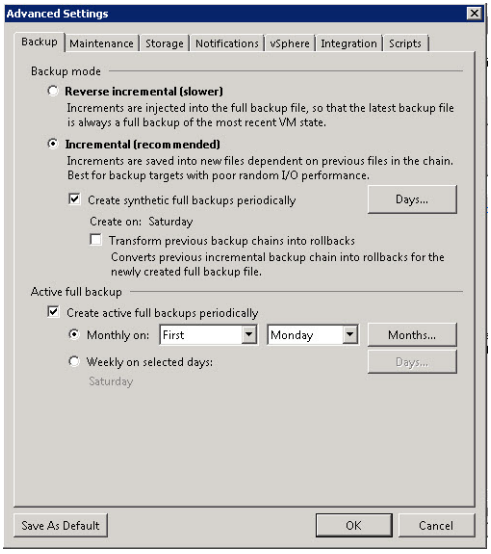

- Ensure **Reversed Incrementals** is unselected (I**ncremental** is selected). Using this option negatively impacts ExaGrid deduplication and will likely cause consumption issue.
- Ensure **Transform previous backups...** is unchecked
- Enable synthetic fulls. See “ExaGrid-Veeam Accelerated Agent” on page 6 for details.
- Schedule **Active full backups** for at least once a month.
- Set other options as needed.

18. In the **Maintenance** tab:

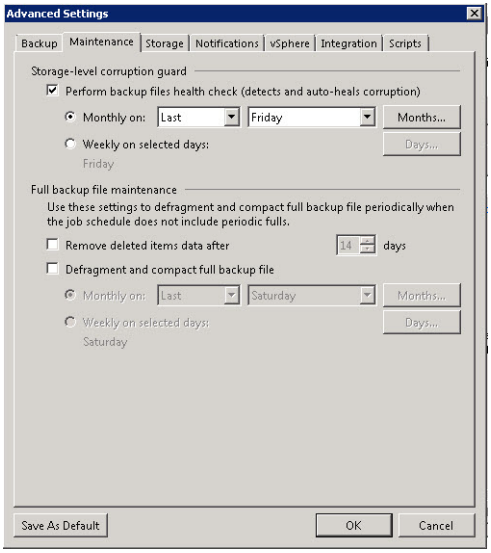

- Ensure that **Perform backup files health check** is enabled on some periodic basis.
- There is no need to do any **Full backup file maintenance**
- Set other options as needed.

19. In the Storage tab:

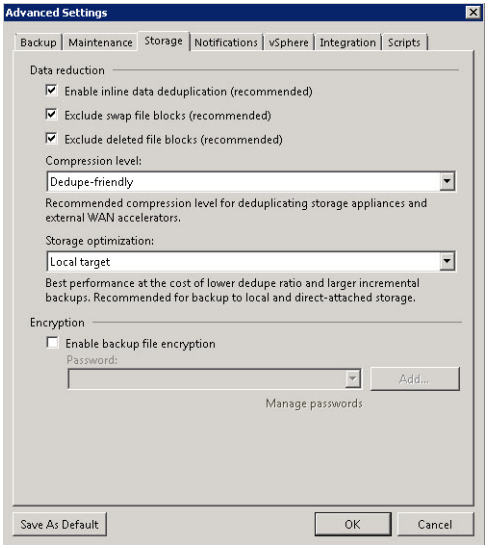

- Check **Enable inline data deduplication.**
- Set **Compression level** to **Dedupe-friendly.**

Caution – The two settings above are critical to achieving maximum overall deduplication and your ExaGrid solution has been sized assuming these settings. Deviating from this practice will likely result in ExaGrid capacity issues.

- Set **Storage optimization** to **Local Target**.
- Ensure **Encryption** is not enabled
- Set other options as needed.

20. In the Notifications tab, set as needed.

21. In the vSphere tab, set as needed.

22. In the Integration tab, set as needed.

23. Click **OK** and the Advanced settings will be applied and the Advanced dialog box will close.

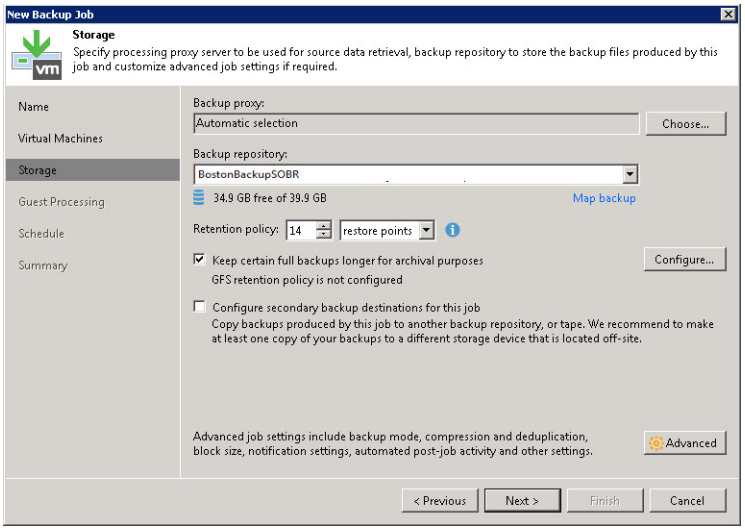

24. Click **Next** and you will be prompted for Guest processing.

25. Make selections as needed.

26. Click **Next** and you will be prompted to schedule the backup:

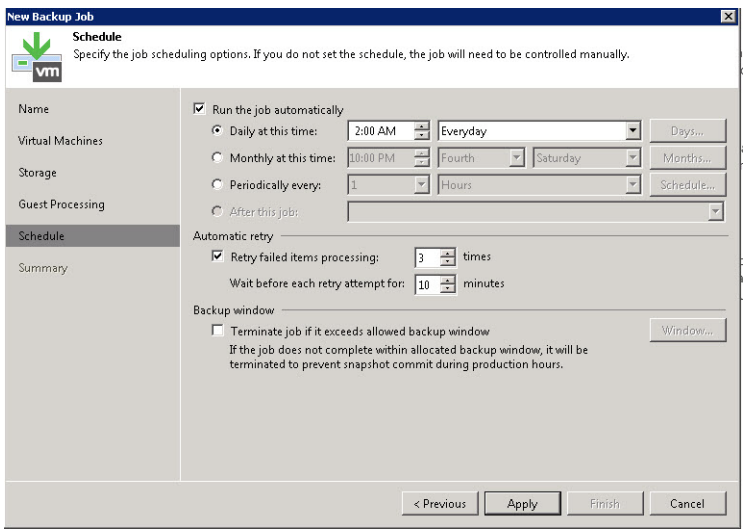

27. Set the schedule as needed.

28. Under **Automatic retry**, do not set **Wait before each retry attempt** to greater than fifteen (15) minutes.

29. Click **Apply** and you will be prompted to review your settings:

30. Click **Finish** and the job will be listed under **Jobs-->Backup** and will run according to your settings:

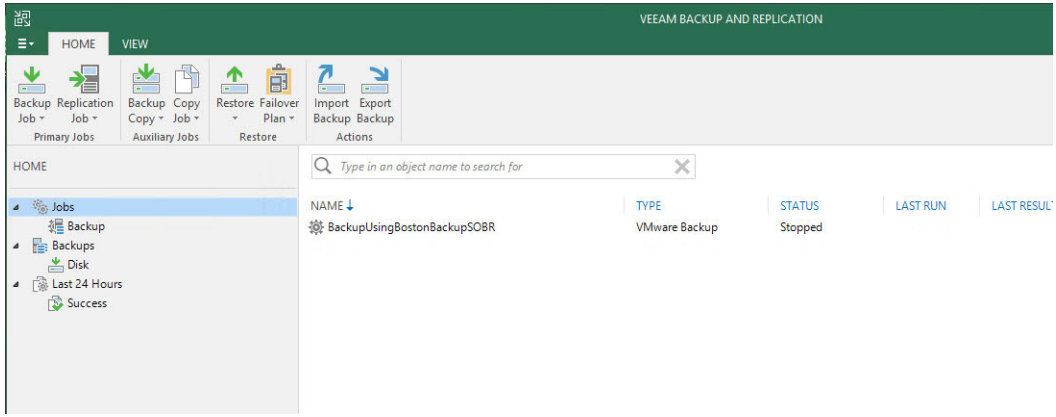

Creating Backup Jobs for Physical Devices and Applications: Step-by-Step

In addition to backing up VMs, Veeam can protect a wide variety of physical devices and applications. Except where explicitly called out, use the same ExaGrid recommended settings for all physical devices and applications.

Veeam's agents for physical servers can be used with ExaGrid as follows:

Veeam Feature	ExaGrid Support
NAS File Shares	ExaGrid Veeam Data Mover repository can be a backup target. First backup is a full, with forever incrementals after that. All best practices described in this document apply.
Agent for Windows	ExaGrid Veeam Data Mover repository can be backup target and includes scheduled synthetic fulls that are accelerated by the Data Mover.
Agent For Linux	If you manage your physical server with Veeam Backup and restore (server mode) then Synthetic fulls and active fulls are available options. These options are not available in Workstation mode. In addition, you can point Linux physical server jobs to an ExaGrid Veeam Data mover repository.
Veeam Plug-ins for Enterprise Applications	ExaGrid Veeam Data Mover repositories can be backup targets for Veeam Plug-ins.
Veeam Backup for Microsoft Office 365	Not supported. Veeam does not support using a deduplication appliance as a target repository for Office 365 backups.

Before You Begin

Before you begin backing up a physical device or application with the Veeam Agent you must create one or more:

- ExaGrid Veeam Shares. See “Creating a Veeam Share” on page 14.
- Veeam Repositories. See “Veeam Repositories” on page 16.
- For some physical devices: Veeam Protection Groups. See“Creating a Veeam Protection Group” on page 47. Consult Veeam documentation for specifics of physical devices.
- For NAS File Shares, create the appropriate Veeam Backup Job using the same ExaGrid SOBR as is used for VM and other backups. See “Veeam Backup Jobs - Overview” on page 30 for job settings.

Creating a NAS File Share Backup Job

To create a NAS File Share Backup Job:

1. In Veeam’s top main menu, click the **Home** tab.
2. Click **Backup Job**.

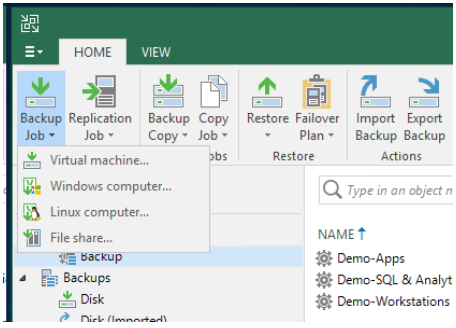

3. From the drop down menu select **File Share...** and the Name dialog box opens:

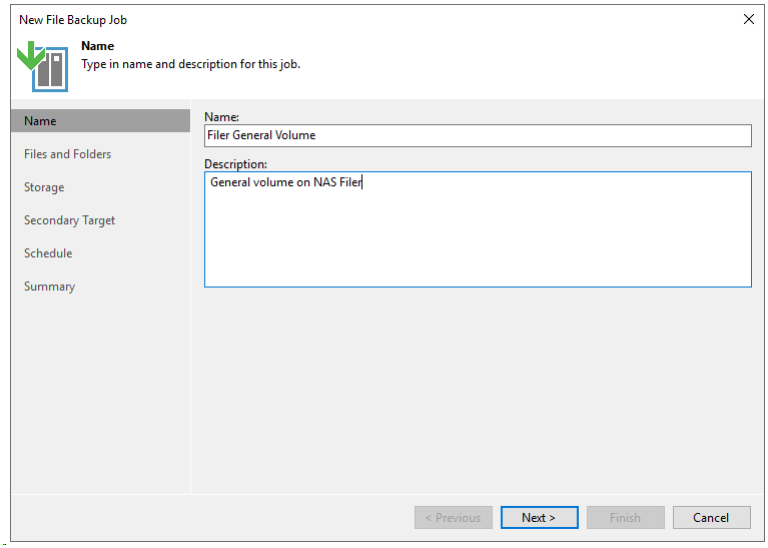

4. In the spaces provided, enter a name and description of the backup job.
5. Click **Next** and the Files and Folders dialog box opens:

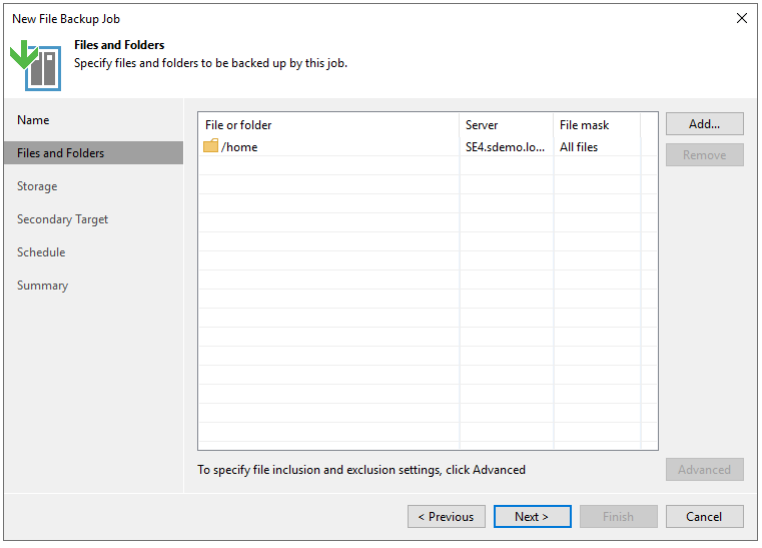

6. Click **Add** and consult Veeam documentation for adding files and folders from file shares to the backup job.

7. When done adding Files and Folders to this Backup Job, click **Next** and the Storage dialog box opens:

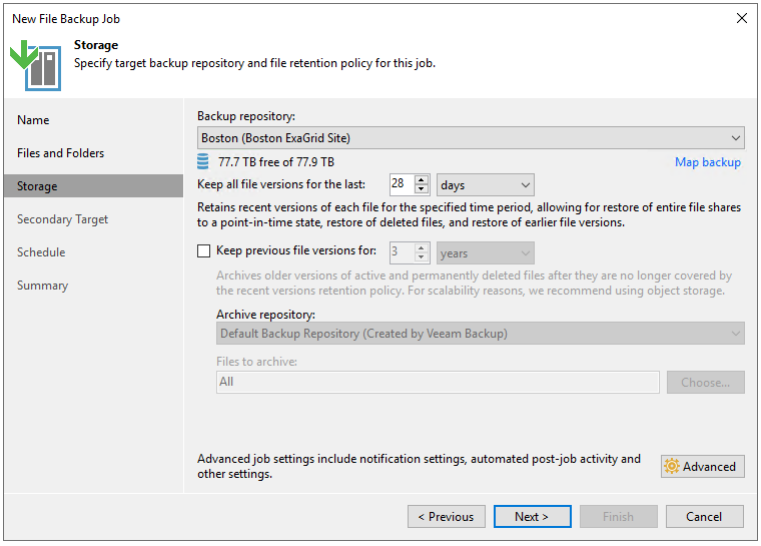

8. From the Backup repository drop down list, select the repository with ExaGrid shares/extents that you created for this backup job. When using a Veeam SOBR, it can receive all Veeam backups - VM, File Share, physical devices, etc.

9. Set the retention of file versions as needed and based on the retention specified when the ExaGrid solution was sized.

10. Previous file versions should be archived on non-ExaGrid storage since the target for the Archive repository is recommended to be an object store and because it cannot be a SOBR.

11. Click the **Advanced** button and the **Advanced Settings** dialog is displayed. ExaGrid specific settings are on the **Storage** tab:

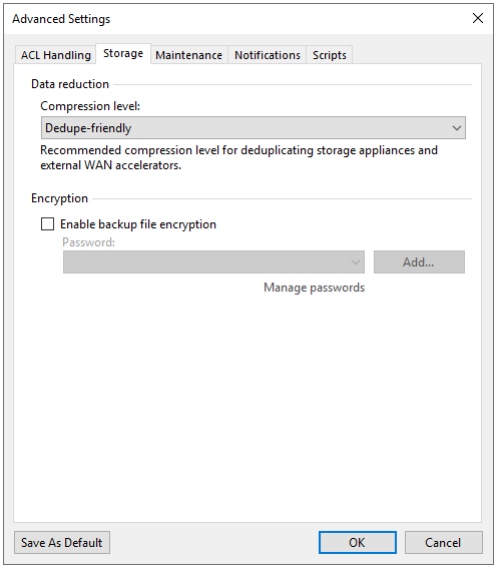

12. Set **Compression level** to **Dedupe-friendly.**

Caution – This setting is critical to achieving maximum overall deduplication and your ExaGrid solution has been sized assuming these settings. Deviating from this practice will likely result in ExaGrid capacity issues.

---

13. Click **Next** and the Schedule dialog box opens:

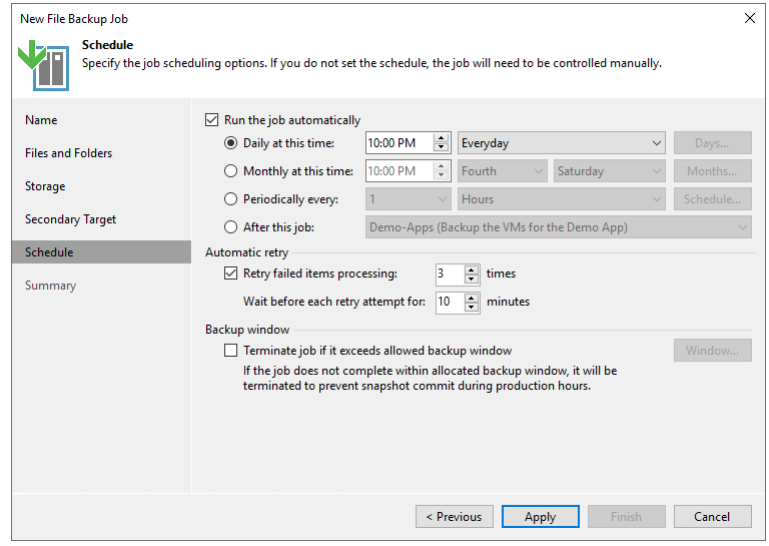

14. Set the schedule as needed.

15. Under **Automatic retry**, do not set **Wait before each retry attempt** to greater than fifteen (15) minutes.

16. Click **Next** and the Summary page is displayed.

17. Click **Finish** and the job will run according to schedule.

Creating a Veeam Protection Group

Protection groups are user created groups of physical servers that allow you to run multiple physical agent backups from a single backup job. In addition, protection groups allow you to automatically manage Veeam agent installations and updates.

To create a Protection Group:

1. From Veeam's left-hand navigation, click **Inventory**

2. From Veeam's left-hand navigation, click **Physical Infrastructure** and the Protection Group Options are displayed:

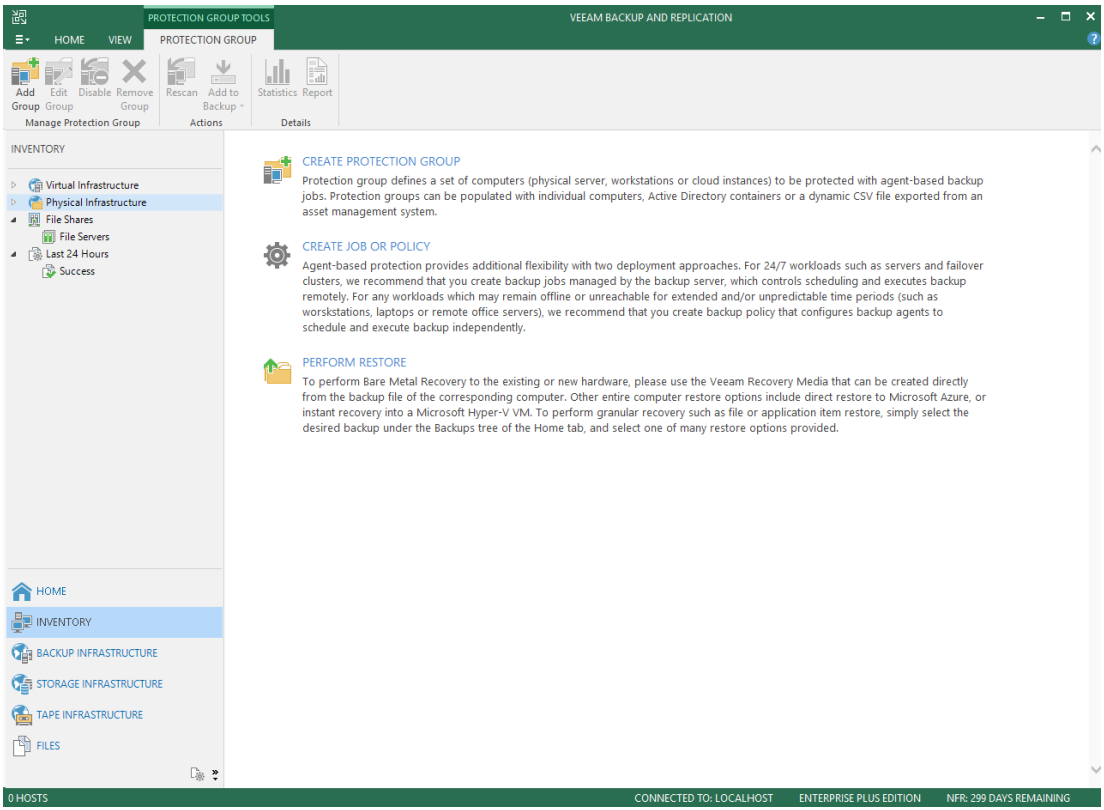

Click **Create Protection Group** and the Create New Protection Group's Name dialog box opens:

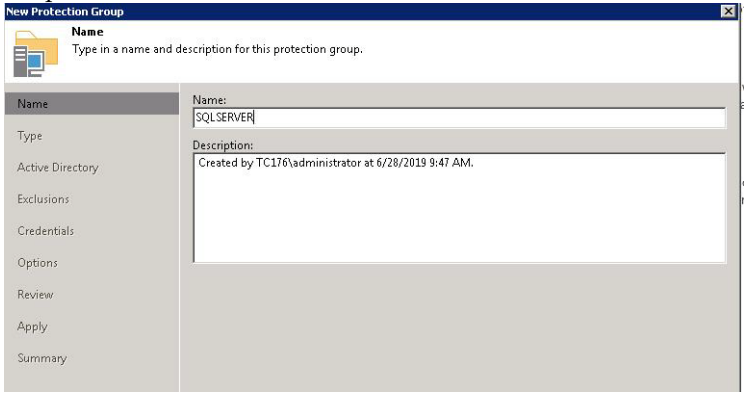

3. In the spaces provided enter a name and description of the protection group.

4. Click **Next** and Type dialog box opens:

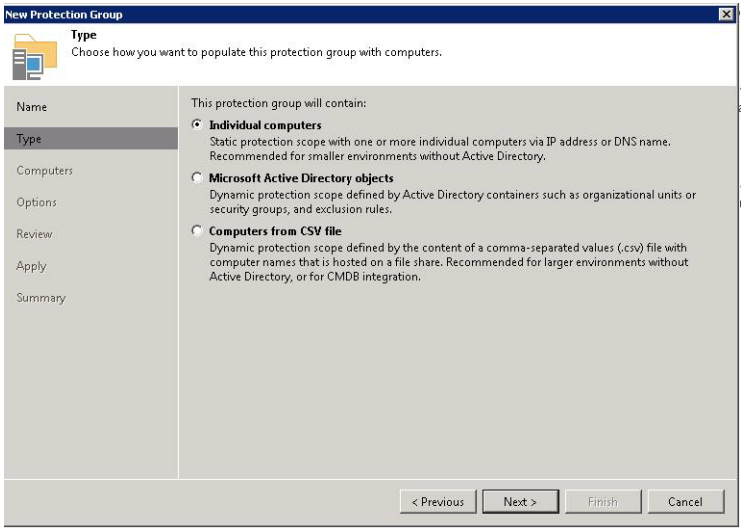

5. Select **Individual Computers**.

6. Click Next and the Computers dialog box opens.

7. Click Add and the Add Computers dialog box opens:

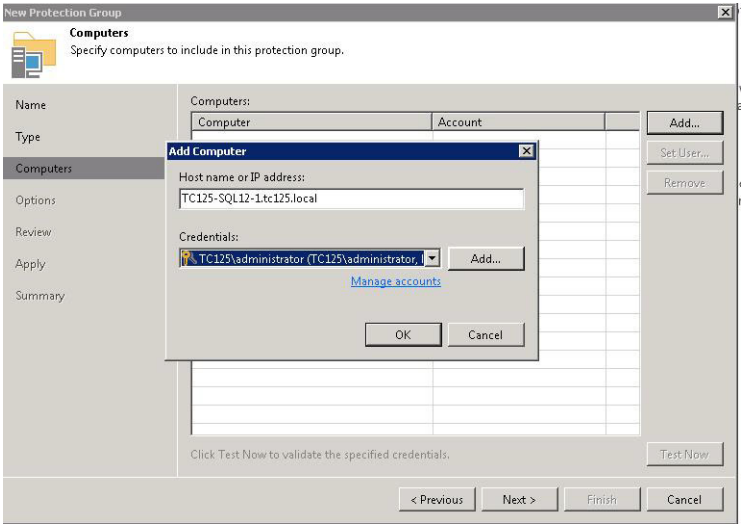

8. In the space provided enter the name of the computer to be added or its IP address

9. Under **Credentials**, either select credentials from the drop down list or add new credentials by clicking **Add**.

10. Repeat the process for each computer you want to add to the Protection Group.

11. Click **Next** and the Options dialog box opens:

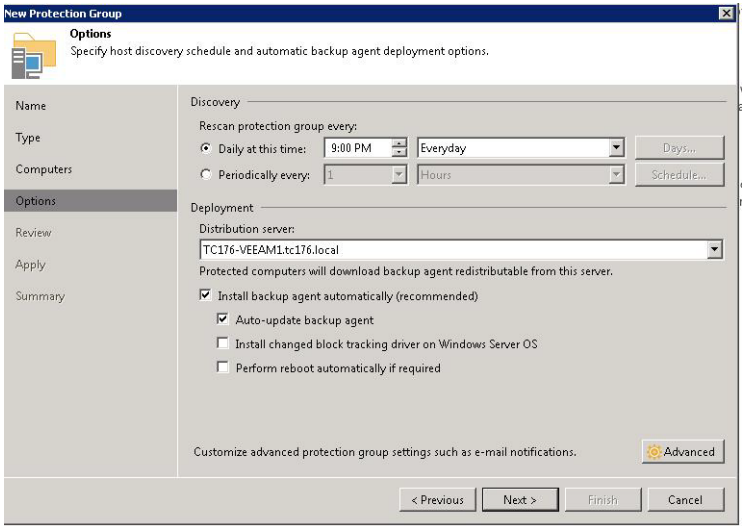

Set the Rescan option in such a way as you know any new computers added to this group will be scanned before the next backup is run.

12. From the **Distribution Server** drop down menu select the Veeam server from which the Veeam Agent will be downloaded.

13. Select **Install backup agent automatically**.

14. Select **Auto update backup agent.**

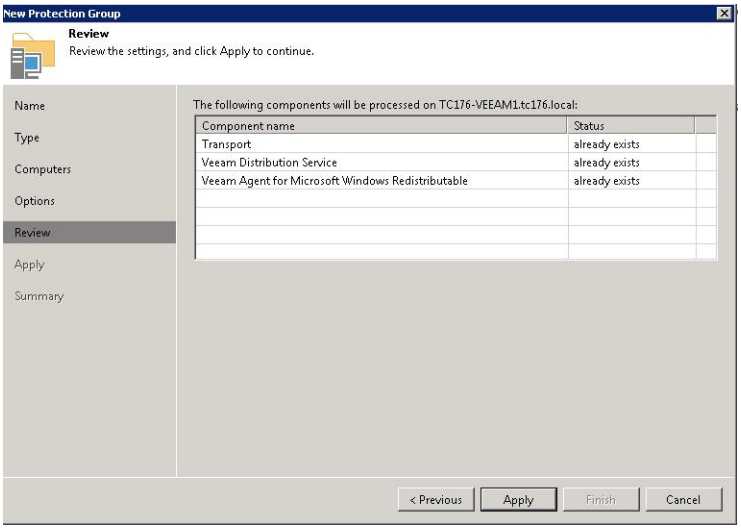

15. Click **Apply** and the Apply dialog box opens:

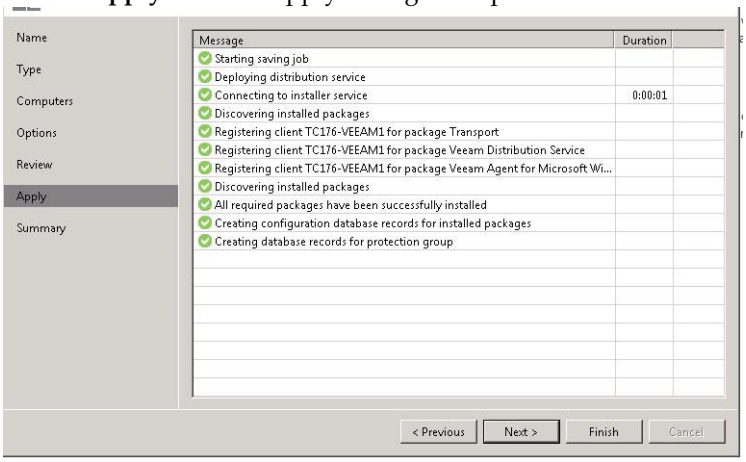

16. From the Apply dialog box you can view the Protection Group's.

17. When the Apply operation is complete, the Protection Group is created and you can either click **Summary** to review the Protection Group's settings or click **Finish** and the operation is complete.

Creating a Physical Device Backup Job

To create a physical device backup job:

1. In Veeam's top main menu, click the **Home** tab.

2. Click **Backup Job**.

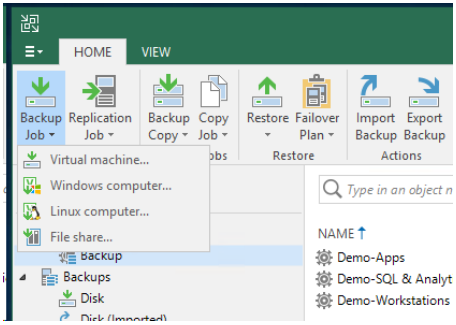

3. From the drop down menu select either **Windows Computer...** or **Linux Computer...** and the Job Mode dialog box opens:

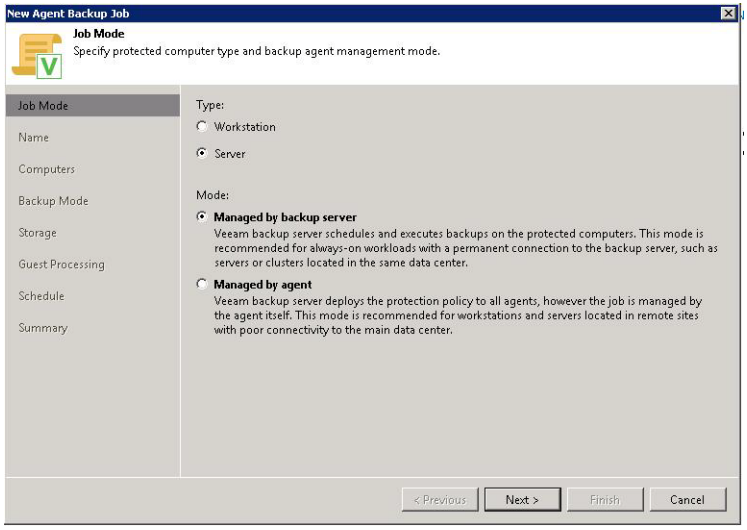

4. Select **Server** mode. Doing so allows you to perform synthetic full backups and active full backups on Windows and Linux physical servers.

5. Click **Next** and the Name dialog box opens:

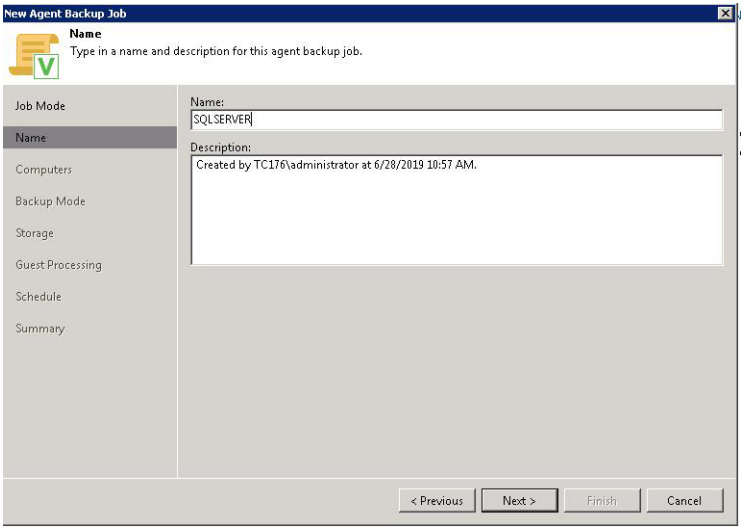

6. In the spaces provided, enter a name and description of the backup job.

7. Click **Next** and the computers dialog box opens:

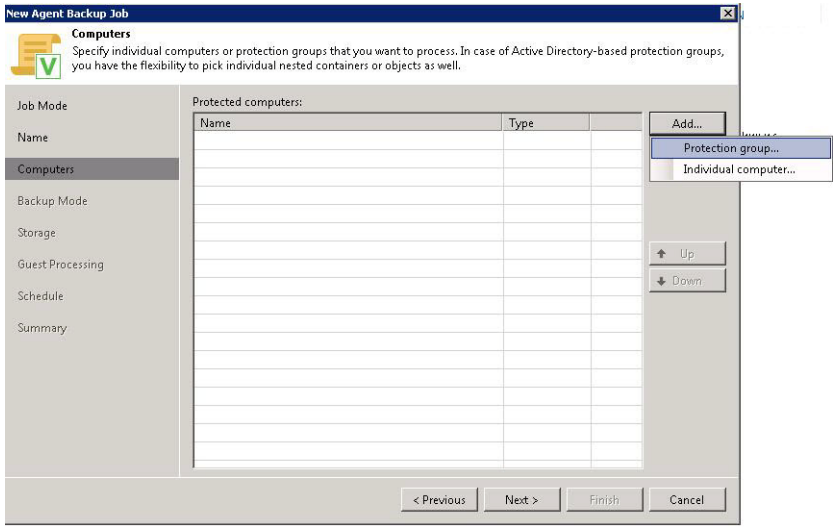

8. Click **Add** and from the drop down menu select **Protection group...** and a dialog box with all available Protection Groups is displayed:

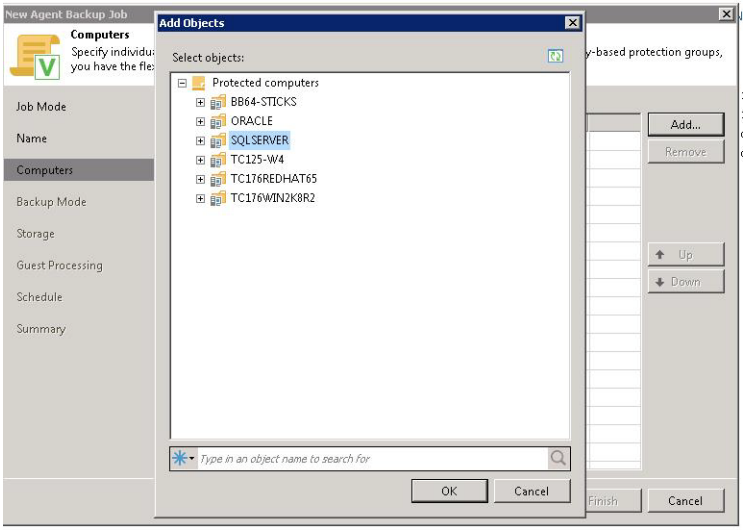

9. Select the Protection Group you created for this backup job.

10. Click **OK** and the group is added to the list of computers.

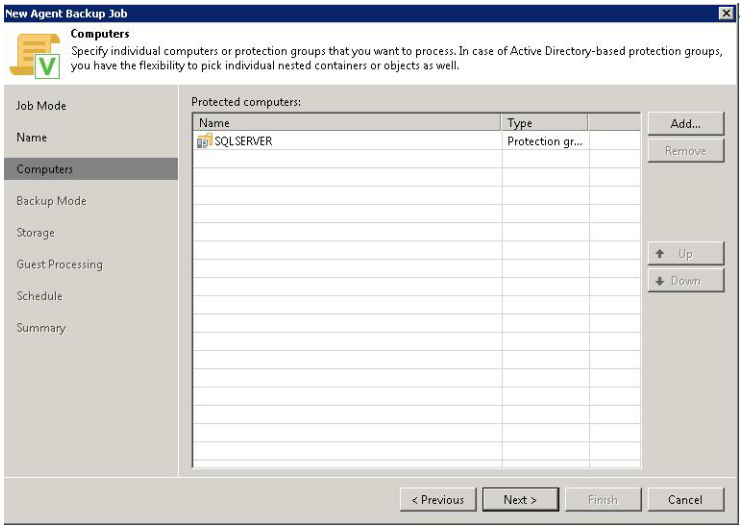

11. Add other groups or computers as needed.

12. When done adding, click **Next** and the Backup Mode dialog box opens:

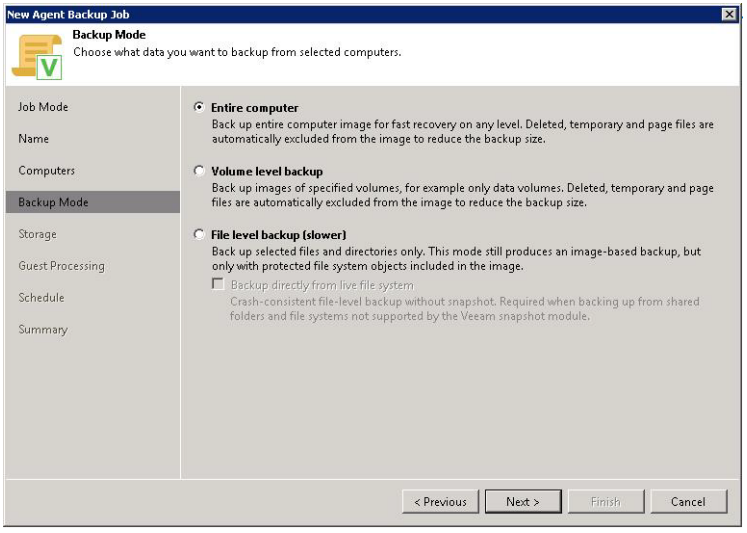

13. Select **Entire computer**.

14. Click **Next** and the Storage dialog box opens:

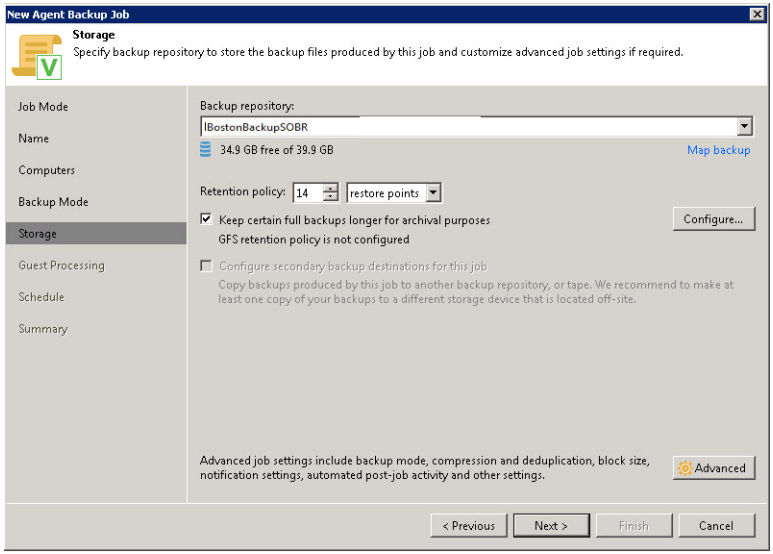

15. From the Backup repository drop down list, select the repository you created for this backup job. When using a Veeam SOBR, it can receive all Veeam backups - VM, File Share, physical devices, etc.

16.

17. Set **Retention Policy Restore points to keep** to no more than 14.

Note – As a best practice, ExaGrid recommends using a **GFS retention policy**. This allows Veeam to retain older backups longer while retaining the fewest number Veeam Retention Policy Restore Points. For example, rather than keep a large number of daily backups, consider retaining 14 or less restore points (days of daily backups) and then with a GFS policy, retain multiple weeks, months, or years of backups.

---

18. Configure a GFS (Grandfather-Father-Son) retention policy:
    - Check the **Keep certain full backups longer for archival purposes**.

- Click **Configure...** and the Configure GFS dialog box opens:

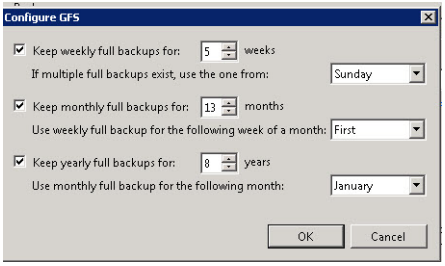

19. Configure the GFS Weekly, Monthly, and Yearly settings. as needed. If your retention requirements have increased beyond the original ExaGrid sizing, please contact your ExaGrid Support Engineer and explain that your retention needs have changed.
    - Click **OK** and the Configure GFS dialog box closes.
20. Check **Configure secondary destination for this job** if, after the job completes, you want a copy of the backup written to tape directly from the ExaGrid landing space,
21. Click **Advanced Settings** and the Advanced Settings dialog box is displayed:

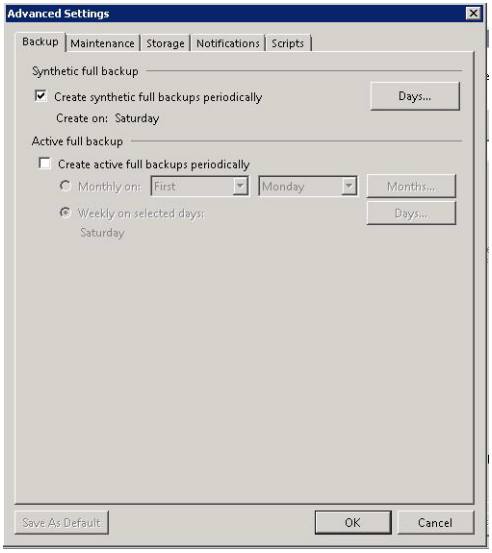

22. If you selected **Server** mode (recommended) in the Job Mode dialog box then you can select the Synthetic full options (recommended).
23. Click **OK** and the Advanced Settings dialog box closes.
24. Click **Next** on the Storage dialog box and the Guest Processing dialog box is displayed

25. Set Guest processing as needed.

26. Click **Next** and the Schedule dialog box opens:

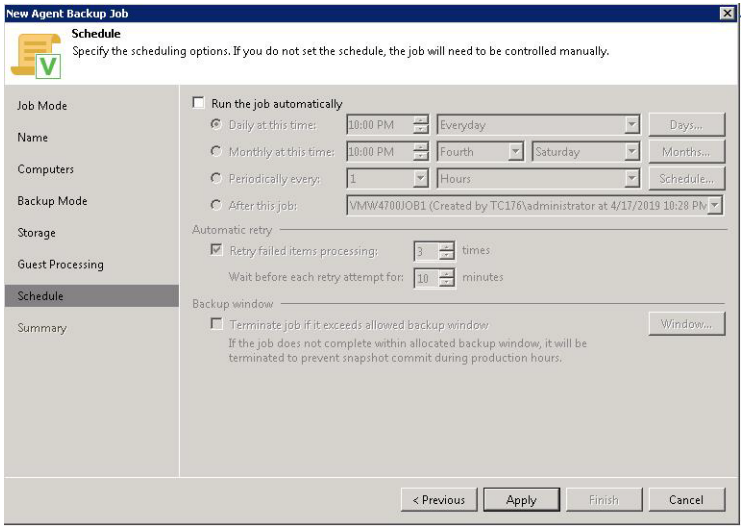

27. Set the schedule as needed.

28. Under **Automatic retry**, do not set **Wait before each retry attempt** to greater than fifteen (15) minutes.

29. Click **Next** and the Summary page is displayed.

30. Click **Finish** and the job will run according to schedule.

Creating a Remote Backup Copy Job

Veeam Backup Copy Jobs are used to copy backup jobs from a local ExaGrid Site to a remote ExaGrid Site.

Backup Copy To a Remote Site Overview

Veeam can copy backups directly from the landing space of an ExaGrid Server in a local ExaGrid Site to the landing space of another ExaGrid Server in a remote (physical) ExaGrid Site. Copy jobs cannot be sent to an ExaGrid VDRT at this time.

Note – It is the ExaGrid-Veeam Accelerated Agent that moves the data from one ExaGrid Site to another. Veeam server resources are not involved. When using Veeam's WAN acceleration, the backup data is moved from the local ExaGrid to the local Veeam server/WAN accelerator, and then from the remote Veeam server/WAN accelerator to the remote ExaGrid.

Copying a backup from a local ExaGrid Server's landing space to a remote ExaGrid Server's landing space provides:

- A Veeam catalog that is "aware" of what is on the local and remote ExaGrid Sites. This ensures recovery from disaster scenarios are as fast and efficient as possible.
- Allows you to perform vPower operations, including Instant VM Recovery operations on both the local source backup and the copy at the remote ExaGrid Site.
- Different retention definitions for the backups stored on the local and remote ExaGrid Sites.

To copy backups to a remote ExaGrid Site, you must create:

- A backup job that uses a SOBR repository that contains all ExaGrid Shares on the local ExaGrid Site.

- A backup Copy Job that uses a SOBR repository that contains all ExaGrid Shares on the remote ExaGrid Site.

Backup Copy To a Remote ExaGrid Site - Step-by-Step

Before you can create a Veeam Backup Copy job that copies a backup from a local ExaGrid Site to a Remote ExaGrid Site, you must

- Create a Backup job at the local ExaGrid Site that uses a Scale-out Backup Repository (SOBR). This SOBR will be used as the source for the Backup Copy job.
- Create a Backup Copy job at the local ExaGrid Site that uses a SOBR on a remote ExaGrid Site.
- For details, on creating repositories see:
  - "Creating Backup Repositories" on page 17.
  - "Creating Scale-out Backup Repositories" on page 25.
- In addition, you must familiarize yourself with all of the concepts provided earlier in this guide.

The following example uses a Veeam Scale-out Repository to copy backups from a local ExaGrid Site in Boston, to a remote ExaGrid Site in Marlborough.

To create a Veeam Backup Copy job:

1. From the **Home** tab in the Veeam user interface, click **Backup Copy Job**.

2. From the drop down menu select **Virtual Machine** or **Physical Machine** and the New Backup Copy Job dialog box opens:

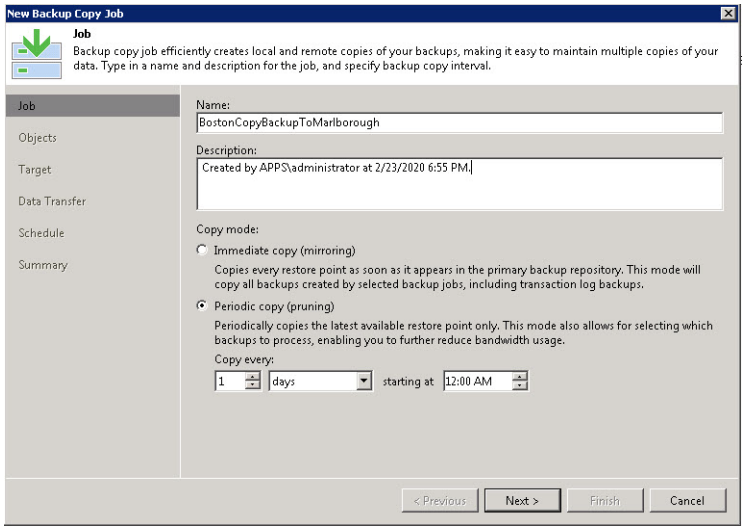

3. In the spaces provided, enter a name and description. Adapt a naming convention that easily identifies the backup copy job.

4. Under **Copy every**:, set the frequency at which you want the copy job to run. For example, set this in such a way that up to 14 days of the most recent backups are kept only on the primary ExaGrid Site/SOBR, and the later backups are copied by the backup copy job to a remote ExaGrid Site/SOBR.

5. Click **Next** and the Objects options dialog box opens.

6. Click **Add** and from the drop down menu select either **From jobs... From Infrastructure...**, or **From Backups...** In our example we have used From jobs... and the From Jobs dialog box is displayed:

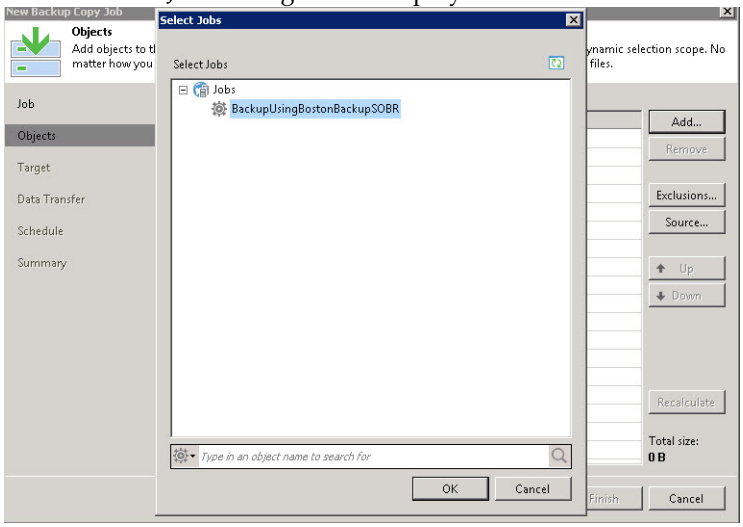

7. Select the backup jobs to be copied.

8. Click **OK** and the Select Jobs dialog box closes and the Object options dialog will be populated with the jobs you selected:

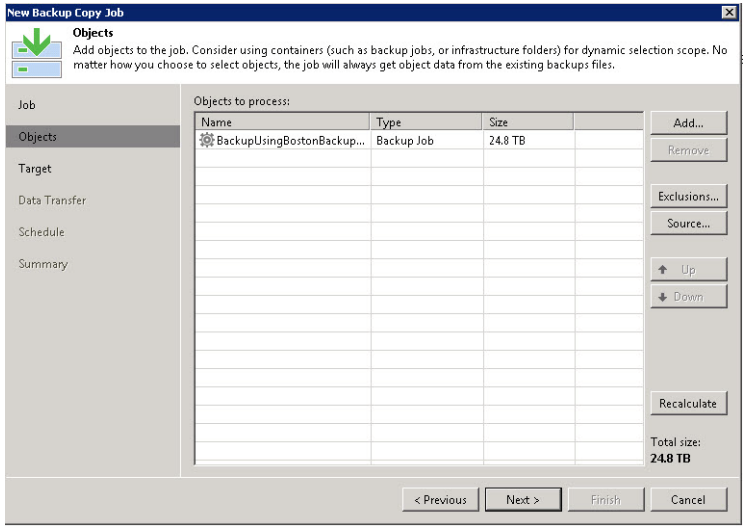

9. Click **Exclusions** to take advantage of Veeam's ability to include or exclude infrastructure folders. This feature allows newly created VMs to be automatically included into the right Veeam backup job. See Veeam documentation for more details.

10. Click **Next** and the Target options dialog box opens:

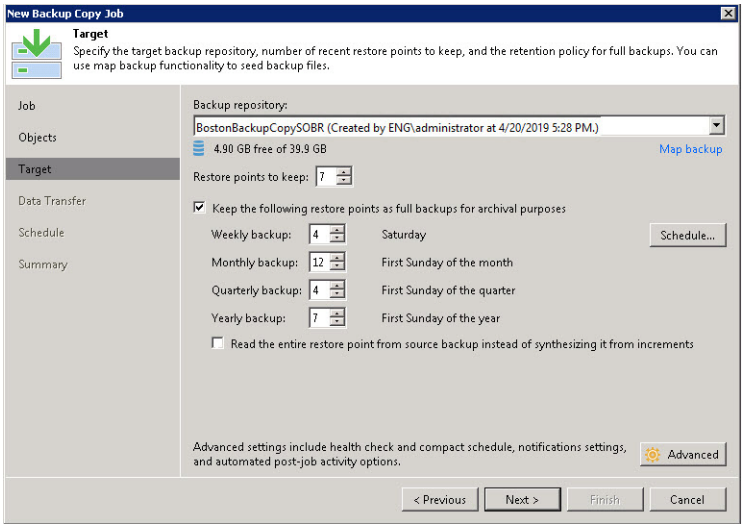

11. From the drop down menu, select the Veeam Backup Copy Job SOBR you created for the remote ExaGrid Site.

12. Set the restore points to be consistent with how your ExaGrid solution was sized by your ExaGrid Systems Engineer. If your retention requirements have increased beyond the original ExaGrid sizing, please contact your ExaGrid Support Engineer and explain that your retention needs have changed.

13. Click **Advanced** and you will be prompted to make advanced backup copy settings.

14. In the **Maintenance** tab:

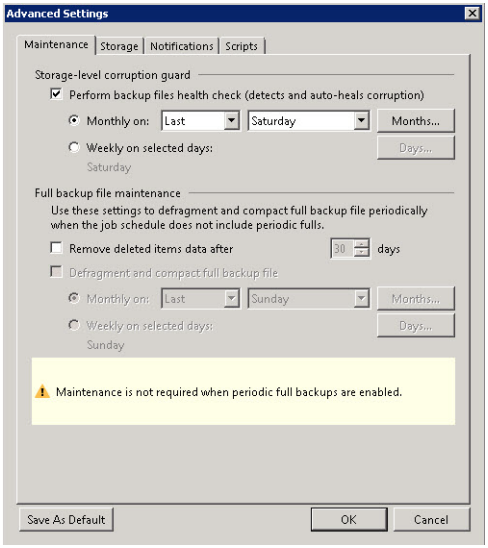

- Ensure that **Perform backup files health check** is enabled on some periodic basis.
- There is no need to do any **Full backup file maintenance**
- Set other options as needed.

15. In the Storage tab:

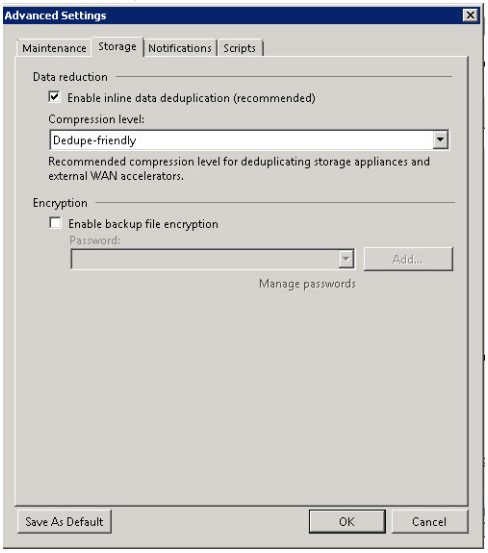

- Check **Enable inline data deduplication.**
- Set **Compression level** to **Dedupe-friendly.**

Caution – The two settings above are critical to achieving maximum overall deduplication and your ExaGrid solution has been sized assuming these settings. Deviating from this practice will likely result in ExaGrid capacity issues.

- Ensure **Encryption** is not enabled
- Set other options as needed.

16. Make settings in the other tabs as needed.

17. Click **OK** and the Advanced settings will be applied and the Advanced dialog box will close.

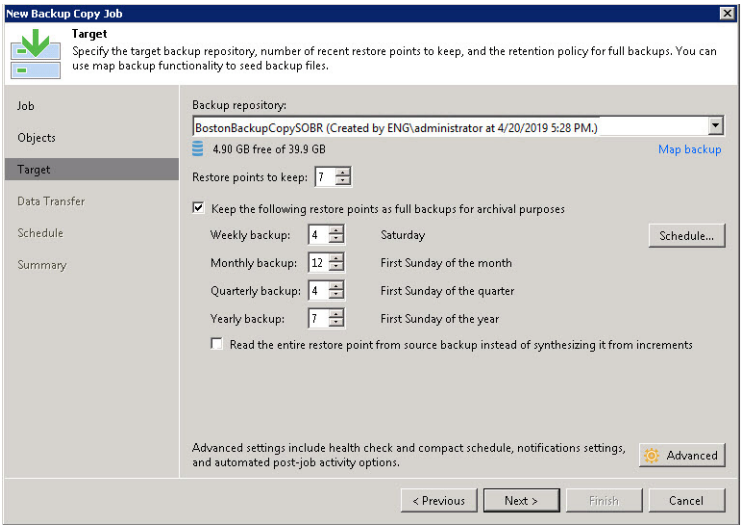

18. Click **Next** and the Data Transfer dialog box opens:

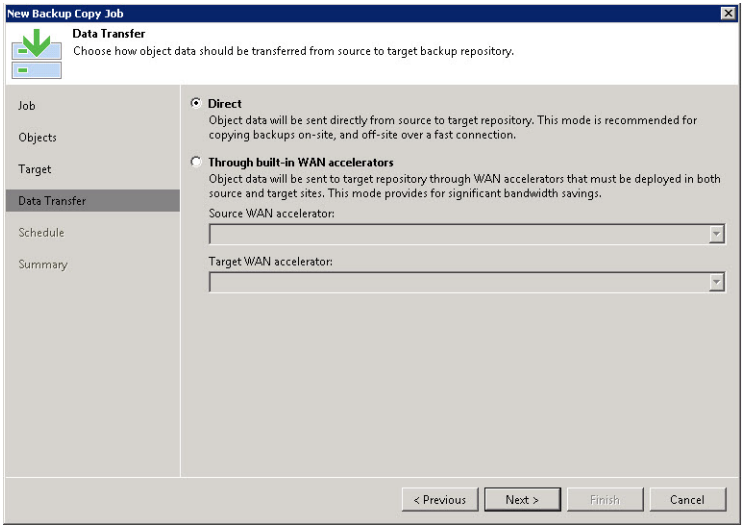

19. Select the data transfer option that best suits your environment.

20. Click **Next** and the Schedule options dialog box is displayed:

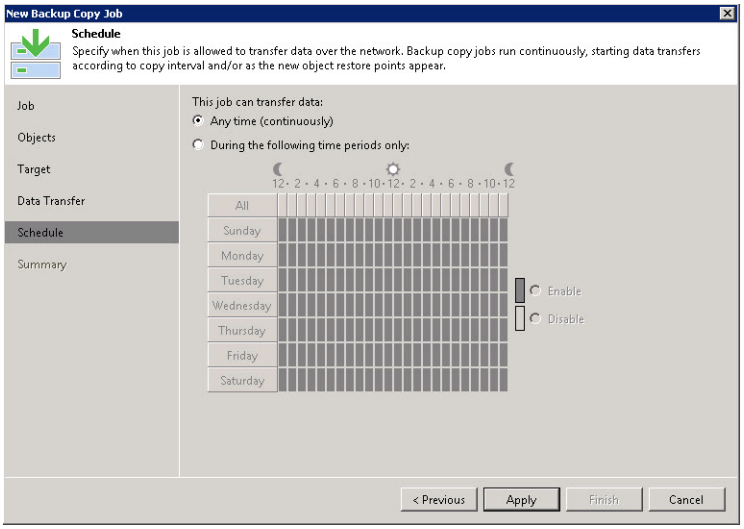

21. Set the schedule as needed.

22. Click **Apply** and you will be prompted to review your settings:

23. Click **Finish** and the job will be listed under **Jobs-->Backup Copy** and will run according to your settings:

24. Select the job and its details are displayed.

Managing Veeam's Scale Out Backup Repositories

Normal operation of Veeam's Scale Out Backup Repository with multiple ExaGrid Servers consists of:

1. The first time a (new) VM is backed up by Veeam using a backup job that targets a SOBR, Veeam chooses the SOBR extent with enough free storage space for that VM's backup. Each SOBR extent is a Veeam Backup Repository that has as its target a Veeam share on a specific ExaGrid Server.

2. All subsequent backups of that VM will be placed on the same extent, since the preferred SOBR policy for ExaGrid is "Data locality". If that extent does not have sufficient storage space, Veeam will attempt the backup, and may or may not succeed depending on the dynamic use of backup storage space on the ExaGrid Server.

Note – Even if the SOBR policy is *Performance*, Veeam will always put the new backups for a VM on the same extent used for previous backups of the VM.

3. The Veeam job Status/Actions and history shows which extent is chosen for a specific VM's backups.

The following sections provide guidance should an issue arise.

Extent (ExaGrid Veeam share) Is Offline

If an extent (e.g. ExaGrid Veeam share on one specific ExaGrid Server) goes offline or becomes unavailable, VM backups going to that extent can either fail, or start a new Veeam backup chain on another extent (another ExaGrid Server) based on the

SOBR's Advanced setting: "Perform full backup when required extent is offline" - see "Creating Scale-out Backup Repositories" on page 25.

Low Backup Storage Space

While Veeam SOBR policies take into account the ExaGrid Servers' free backup storage space, the dynamic changes in backup storage can still occasionally lead to a Veeam backup failing due to lack of ExaGrid Server backup storage free space.

If the backup storage space shortage is a temporary condition, and ExaGrid adaptive deduplication is able to make landing space available in time, Veeam's retry of the backup job can succeed.

Extent Service Actions

Consult Veeam documentation for more details about the service actions that can be performed on SOBR extents - such as maintenance mode, seal mode, evacuation, etc. Contact your ExaGrid customer support engineer before performing any of these service actions as they should not normally be needed.

Converting non-SOBR Configuration to SOBR

Consult with your ExaGrid customer support engineer to develop a plan to move from a non-SOBR configuration to a SOBR configuration.

While you can easily combine existing Veeam repositories into a single Veeam SOBR, having a Veeam SOBR that contains multiple extents (repositories) from the **same** ExaGrid Server can lead to Veeam thinking there is more free backup capacity than actually exists. Avoid this by **not** adding existing Veeam repositories into a Veeam SOBR, and use the following steps:

1. Create a new share on each ExaGrid Server and create a new Veeam SOBR out of them as described in this document.

2. Create or clone Veeam jobs to use the new SOBR. The new backups will likely land on different ExaGrid Servers than before using SOBR, but will still be deduplicated against the older backup jobs.

3. Disable the old, non-SOBR jobs.

4. Veeam will have knowledge about the old non-SOBR jobs so they can be used for restores, however since they are no longer actively running, Veeam will not delete the older backups that are outside retention policies.

5. Work with ExaGrid customer support to use Veeam to periodically delete older backups on disk, removing them from the ExaGrid server. If your retention settings are such that all older backups eventually are removed from the older ExaGrid shares, you can then delete the ExaGrid share.

ExaGrid configurations with a single Veeam share configured on each ExaGrid Server can be easily converted to a SOBR by simple adding each existing Veeam repository into the SOBR. Veeam will prompt to automatically update existing jobs to use the new SOBR.

Expanding Veeam's Scale Out Backup Repository

Note – Some Veeam Editions limit the number of extents (ExaGrid shares) that can be added to a SOBR repository. Consult with your Veeam representative on these limits.

When additional ExaGrid Servers are added to an ExaGrid Site:

1. Create a single ExaGrid-Veeam Accelerated Agent share on the new ExaGrid Site(s). See"Creating a Veeam Share" on page 14.

2. Create two Veeam ExaGrid-type repositories for each new share, one for backup jobs, the other for backup copy jobs. See "Creating Backup Repositories" on page 17.

3. Add the newly-created ExaGrid-type backup job repository as an additional extent to the existing backup job SOBR, and add the newly-created ExaGrid type backup copy job repository as an additional extent to the existing backup copy job SOBR.

Veeam will automatically see the additional free backup storage available through the expanded SOBR repository extents and start to direct per-VM backup files as needed to the new storage.

Note – Because of Veeam's enforcement of the Data Locality policy for all deduplicating storage, only new VM backups will take advantage of the additional free backup storage.

ExaGrid Veeam Share Migration

With Veeam's Scale Out Backup Repository, the need to migrate a Veeam share from one ExaGrid Server to another Server is significantly reduced.

To migrate an ExaGrid Veeam share that is part of a SOBR, use the following steps:

1. Put the Veeam SOBR extent for the ExaGrid Veeam share being migrated into maintenance mode.
2. Remove the Veeam SOBR extent from the SOBR. Do **not** evacuate the extent.
3. Delete the Veeam Backup Repository for the ExaGrid Veeam share. This will not delete backups on the ExaGrid Server.
4. Migrate the ExaGrid Veeam share to a different ExaGrid Server.
5. Create new Veeam Backup Repositories for the new ExaGrid share location.
6. Add the new Veeam Backup Repositories to the Veeam SOBR repositories.
7. Scan the SOBR repository. This should discover the Veeam backup chains on the new extent.

Veeam SOBR Support Resources

Veeam has multiple knowledge base articles available that describe SOBR behaviors/policies, and in some cases, registry key settings used to influence them. To use Veeam SOBR, ExaGrid does not require any special registry key settings. Work with your ExaGrid and Veeam Customer Support engineer to resolve any interoperability issues between ExaGrid and Veeam's Scale Out Backup Repository.

Veeam and ExaGrid Deduplication Ratios

ExaGrid Systems recommends configuring Veeam to deduplicate data within a job. Typically Veeam will achieve a 2:1 deduplication ratio.

Once the deduplicated Veeam backup data has landed, the ExaGrid System will further deduplicate the Veeam data anywhere from 3:1 to 7:1.

The ExaGrid System will report how well it deduplicated the Veeam backup data. The ExaGrid System will not report how well Veeam deduplicated the data.

To calculate the overall deduplication ratio achieved by both Veeam and your ExaGrid System; multiply the typical Veeam deduplication ratio of 2:1 by the deduplication ratio reported by your ExaGrid System. For example:

```
Veeam 2:1 x ExaGrid 7:1 = 14:1
```

To determine the exact deduplication achieved by your Veeam/ExaGrid System solution, use the Veeam Deduplication Report PowerShell Script.

Using the Veeam Deduplication Report PowerShell Script

The Veeam PowerShell Script generates a report that details the deduplication ratios achieved by your Veeam server. The ratios in the PowerShell report will calculate the overall deduplication ratios achieved by your ExaGrid/Veeam backup solution.

To use the Veeam PowerShell Script

- Ensure that Veeam's PowerShell Snapin is installed on your Veeam Server. The PowerShell Snapin is installed from the Veeam Backup & Replication Installer as described in: http://www.veeam.com/kb1489#/kb1489
- Download the script to your Veeam server:

1. From the ExaGrid main menu, click **Help**.
2. From the drop down menu, click the **Online Library**. The on line library is displayed.
3. Open the **Named VM Backup Applications** section.
4. Right click on the **Veeam Deduplication Report PowerShell script**
5. **S**elect the menu option to save the script as **`VeeamDedupReport.ps1`** to your Veeam server.
6. From the drop down menu in the upper left-hand corner of the Veeam user interface, select **Console**.

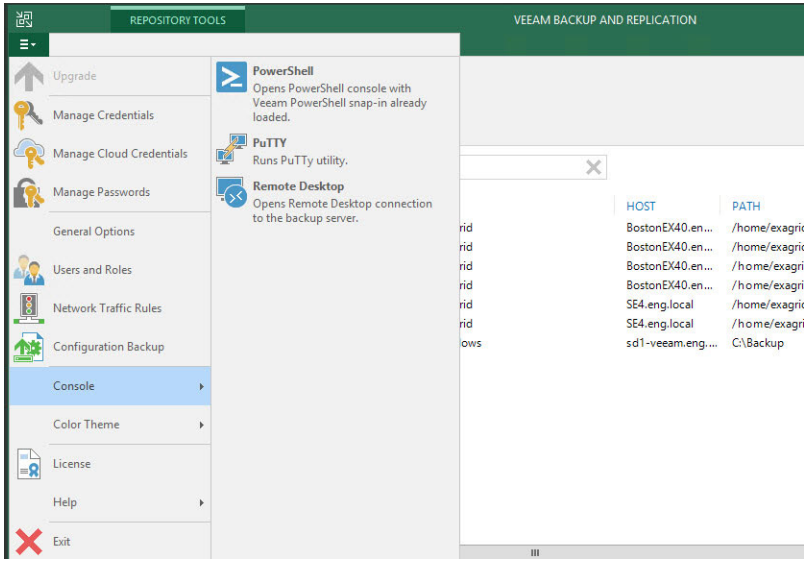

7. From the secondary menu, select **Power Shell** and the Veeam PowerShell window opens.
8. From the Power Shell window, navigate to the script's location.
9. At the prompt, enter:

   `.\VeeamDedupReport.ps1 -egridDedup` <ExaGrid Ratio>

Only enter the antecedent portion of the ExaGrid deduplication ratio i.e. 7.5 **not** 4.5:1. The -egridDedup argument is optional.

For example, if your ExaGrid System is reporting 7.5:1 deduplication ratios for your Veeam backups, then:

```
.\VeeamDedupReport.ps1 -egridDedup 7.5
```

Interpreting the Veeam PowerShell Results

The PowerShell Script report is organized by Veeam repository and includes:

- Job name
- Combined Veeam and ExaGrid deduplication ratio (antecedent only). The Combined ratio is calculated using the ExaGrid deduplication ratio multiplied by the Veeam deduplication ratio.
- Veeam only deduplication ratio (antecedent only). The Veeam deduplication ratio is calculated by dividing the Total Backup Size by the Total Data Size.
- Number of backups
- Compression setting
- Whether deduplication enabled
- Imported and deleted jobs are not included in this report.

Jobs that were deleted without deleting their restore points are still displayed but without compression and deduplication results.

Tape Jobs

ExaGrid's unique landing space architecture allows you to take advantage of **Veeam's Tape Job** feature. With Tape Jobs you can you can write backups to an ExaGrid System, then write those backups directly from the ExaGrid landing space off to tape.

This section provides a high-level description of using Veeam's Tape Job feature with an ExaGrid System. Consult Veeam's documentation for detailed instructions and best practices for using the Tape Job feature.

To use Veeam's Tape Job feature with an existing Veeam backup job:

1. From Veeam's user interface, create a tape job that uses the existing backup job that points to an ExaGrid Share.
2. Configure the tape job to run as soon as the backup job completes.
3. Follow Veeam's the best practices for creating tape jobs.

Recovering From the ExaGrid Source Share

Recovers are done from the Veeam user interface. Consult the Veeam documentation for details.

Disaster Recovery

In the event of a disaster, you can recover Veeam backups that your ExaGrid System has replicated to an ExaGrid InstantDR Share at a remote location.

To recover a Veeam backup from an ExaGrid InstantDR Share:

1. Create a Veeam repository for every ExaGrid InstantDR Share.
2. Assemble the repositories into a SOBR repository.
3. Right click on the newly created SOBR repository and click **Rescan repository**:

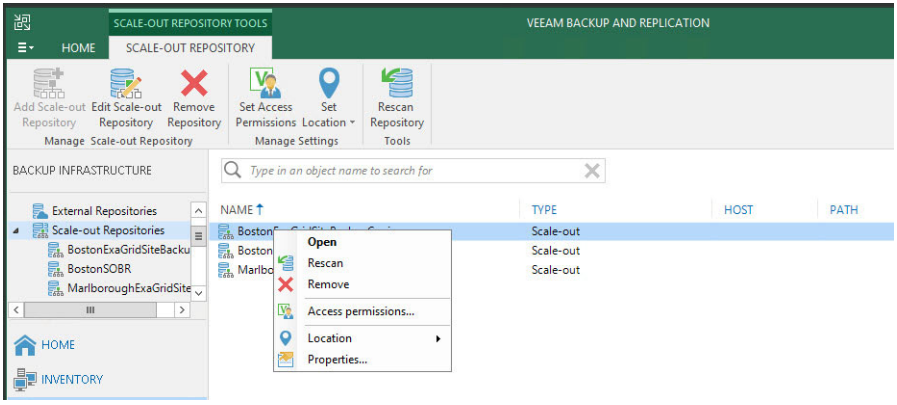

4. When the Rescan completes, in Veeam's left-hand navigation frame click **Backup & Replication**.
5. In the Backup & Replication tree, under **Backups**, click **Imported.**

6. Right click on the imported (rescanned) backup and perform the restore operation needed:

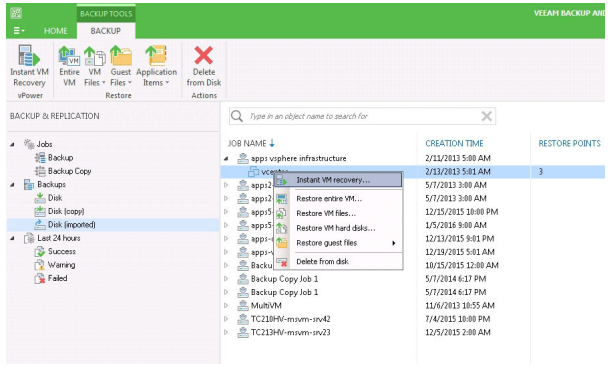

Note – You cannot use Veeam's Instant recovery on a backup imported from an InstantDR Share; you can use Veeam to do a full VM restore, file level recovery, or application item recovery from an InstantDR Share.

In case of a disaster, if a Veeam Share is replicated (using ExaGrid) to another ExaGrid Site, the Veeam share can be recovered to the ExaGrid Site that hosts the replicated share. The recovered ExaGrid Share can then become the source share. In this case, because the recovered share is now the source share, you can take advantage of Veeam's vPower features, including Instant VM Recovery.

Using Veeam Backup and Replication™ Software with an ExaGrid System

Part Number: 210-0500-03