




ExaGrid Appliances: Backup targets for every need

By Logan G. Harbaugh, senior analyst and validation engineer
 SSG-NOW
 July 2015

Backups are a necessity, and as mixed media, data mining and the Internet of Things drive data usage ever higher, the need to fit more data into small backup windows is becoming more critical for many organizations. One simple solution to this issue is to improve the performance of the backup targets, with systems such as the ExaGrid's EX series of appliances, which can scale out to as many as 25 EX32000E appliances in a scale-out grid, for a combined usable capacity of 1575TB, a total ingest rate of 187.5 TB/hr., and single-pane administration of the whole grid through a console that supports virtually all the major backup software packages.

The ExaGrid appliance consists of two areas of disk storage, one used as a landing zone and one used for deduplicated storage. The landing zone takes data from a backup, which is then deduplicated and stored on the retention area. Deduplication is sufficient for 16 weeks of full backups at the rated capacity, which is generally half of the usable capacity of the appliance. In our testing, the backups were deduplicated quickly into the retention area allowing for the landing zone to be ready for additional backups. The most recent backups remain in the landing zone in an un-deduplicated form for fast restores and VM boots. As new backups come in, they overwrite the un-deduplicated backups in the landing zone always ensuring that the most recent un-deduplicated backups are ready for fast restores and VM boots. Other market solutions only store deduplicated data, which results in slow restores and VM boots as 90% or more of restores and VM boots come from the most recent backup. Though with only one server running 145 MBps to the appliance, it was nowhere near the maximum rated intake of 7.5TB/hr.



Setup

Setting up shares to back up to is as simple as starting a wizard, with most major backup software packages explicitly supported, and generic shares in CIFS, NFS, supported as well. The options are simple and well explained, with security provided by whitelisting of systems that can back up to the grid, either by explicit IP address, or a range of IPs. In addition to transport via CIFS and NFS, the Symantec Open Storage standard (OST) is supported, as is the ExaGrid-Veeam Accelerated Data Mover transport, which enhances performance substantially (up to 6x) when using the Veeam Synthetic Full backup. There's even a utility setting that turns off all deduplication and compression, if desired. Replication of backups to a secondary site is as simple as making a check in a checkbox, and this can be set up after the initial share is created as well. The system also supports the ability to run VMs from storage that products like Veeam and Symantec NetBackup support.

Documents are available through the help menu that guide the administrator through the subtleties of specific configurations, such as working with each particular backup software package that is supported, or using the OST or Veeam transports. The help documents are clear and detailed, useful for the administrator who needs to know about the special ins and outs of working with any particular backup software and the ExaGrid appliances.

Administration

Creating volumes for specific backup software is a simple, wizard-based process, and multiple copies of a volume can be created by using the first volume defined as a template. ExaGrid supports as many as 25 backup applications and utilities, including Arcserve backup, Symantec Backup Exec, Commvault Simpana, HP Data Protector, IBM Spectrum Protect (formerly Tivoli Storage Manager), StorageCraft ShadowProtect, Symantec Backup Exec, Symantec NetBackup, and EMC NetWorker. Native database backups supported are Oracle RMAN and Microsoft SQL Server Dump. Virtual Machine backup applications are Veeam Backup and Replication and Dell vRanger. For backup applications not specifically supported, a Universal Backups format and a Utility (with no deduplication or replication) is also supported.

Replication for a share can be set up at the time of the initial creation of the share, or at a later date. ExaGrid also supports Veeam's Instant Recovery – with a deduplicated volume, there's no delay to 'rehydrate' the share before the instant recovery can take place.

Access to shares is controlled by whitelist, rather than an access control list or Active Directory or LDAP integration. For each share, access is limited to an IP address or range of addresses, or by hostname.

Restores

In general, restores are easy to use and quick to create, whether of single files or bare-metal restores of servers, VMs or volumes. While this is generally done through the backup software's interface, the ExaGrid system makes the process responsive and fast due to the Landing Zone architecture. ExaGrid works with Veeam's Instant Restore feature to run a backed up VM directly from the Landing Zone on the appliance. This makes the VM available in less than a minute, and it can then be moved to a VMware server using the VMware vRealize Suite (formerly vCenter).

Plenty of Room for Growth

While many backup systems can add a shelf or two, the ExaGrid system could start with a single EX1000 with 1TB of capacity and grow that to hundreds of terabytes by adding any one of ten various-sized appliances with suitable capacity – all the systems are compatible, can be managed through the same interface and can be part of the same grid. Adding

another appliance to a grid is as simple as generating a key with the new appliance, and entering the new appliance's IP address and the key through the administration console for the grid. Capacity for the grid as a whole is immediately available, and shares defined on one host can be transparently migrated to any other host.

SSG-NOW's Take

ExaGrid's appliances offer all the features an administrator might look for in a backup target, with great ingest speeds and scalability. ExaGrid offers ten appliances, from 1TB to 32TB, with scale-out grids of up to 25 appliances, allows for an inexpensive entry point and enough capacity for even large data centers. Since speed grows with the number of appliances, and all appliances from the smallest to the largest, run the same software, there's no need to buy more storage than necessary to get all the features you want. This keeps costs down, since dollars per gigabyte is always going down with advances in capacity. •

About Storage Strategies NOW™

Storage Strategies NOW™ (SSG-NOW) is an industry analyst firm focused on storage, server, cloud and virtualization technologies. Our goal is to convey the business value of adopting these technologies to corporate stakeholders in a concise and easy-to-understand manner.

Note: The information and recommendations made by Storage Strategies NOW are based upon public information and sources and may also include personal opinions both of Storage Strategies NOW and others, all of which we believe to be accurate and reliable. As market conditions change however, and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. Storage Strategies NOW, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document.