ExaGrid vs. Dell EMC Data Domain

Comparison of the Approaches and Benefits

FUNCTIONALITY	DELL EMC DATA DOMAIN APPROACH	DELL EMC DATA DOMAIN CHALLENGES	EXAGRID ARCHITECTURAL APPROACH / BENEFITS
SUPPORTS MARKET LEADING BACKUP APPLICATIONS	Yes		Yes
LARGEST FULL BACKUP IN A SINGLE SYSTEM	1.5 petabyte	Maximum ingest equates to a 13-hour backup time. Forces forklift upgrades for performance reasons.	2.69 petabyte
LARGEST FULL BACKUP ENVIRONMENT	No limit - deploy multiple systems		No limit - deploy multiple systems
AGGRESSIVE DATA DEDUPLICATION	8KB blocks with variable length content splitting. Uses 1/20th of the disk with deduplication (20:1) at 18 copies of retention.	Block level deduplication forces architecture of scale-up storage and deduplicated pools of storage for restores/ VM boots.	Uses zone stamps with byte level compare and uses 1/20th of the disk with deduplication (20:1) at 18 copies of retention.
GLOBAL DEDUPLICATION WITH PERFORMANCE LOAD BALANCING	Supports global deduplication in a single controller scale-up system.	Improved deduplication rate but is still bound by a single throughput front-end controller.	Supports global deduplication in a single scale-out system across full resource appliances. Allows flexible backup job management to be configured for performance load balancing as deduplication occurs within and across all appliances.
STORAGE ARCHITECTURE	Scale-up – fixed front-end controller with disk shelf capacity add	As data grows, compute/memory/network ports are fixed. The backup window grows to a point where a larger/faster front-end controller is required. This is called a "forklift upgrade" and is costly and disruptive.	Scale-out: full appliances with all resources - processor, memory, network ports, and disk - are added as data grows, which eliminates costly and disruptive forklift upgrades.
LONG-TERM COSTS	Forklift upgrades	Front-end controllers can hit an ingest limit before the full disk capacity is added, forcing forklift upgrades more often.	Full resource appliances are added in into a scale-out system for a pay-as-you-grow model without any forklift upgrades.
	Maintenance and support is a percentage of list price	Maintenance and support is charged as a percentage of list regardless of what is actually paid for the systems.	Maintenance and support is a percentage of the price paid for the system versus list. M&S will not increase more than 3% per year for the life of the use of the appliances.
	No protection in pricing model	Initial entry price is often lower to gain data center foot-print, and future prices are higher.	The price paid for the appliances is the price protected for 5 years - no increases.
	Constant obsolescence	In order to stay current with the latest processor technology, controllers are obsoleted on a regular basis, and astronomical maintenance and support charges are imputed.	All appliances are supported as long as you run them. Appliances of any age/generation and size work together in a scale-out system for long-term investment protection.
	Hidden costs	Dell EMC charges additionally for replication, DD Boost to increase ingest to a usable level, training, installation, and other.	No hidden costs: replication, full customer support, point and full version releases, installation, training, all failed hardware replacement, and health status monitoring are all included.
INGEST PERFORMANCE FOR THE SHORTEST BACKUP WINDOW	Deduplicates data on the way to disk, called "inline" data deduplication.	Deduplication is extremely compute-intensive. The fastest Data Domain can ingest at 41TB/hr. As a result, Dell EMC has to put software (DDBoost) on the media servers and database servers to leverage more compute. They also need to add flash/SSD to accelerate their look-up tables. DD Boost takes compute from other processes, and flash/SSD increases price. With DDBoost and flash/ SSD, ingest is still only 94TB/hr. for a 1.25 petabyte full backup. This ingest can only keep up with a portion of the rated capacity.	Backups are written directly to disk avoiding the high compute overhead of data deduplication. Deduplication occurs adaptively during the backup window but not inline, delivering ingest performance of 432TB for a 2PB full backup (which is 3 times that of Data Domain) and resulting in a shorter backup window.
RESTORE, VM BOOT, AND OFFSITE TAPE COPY PERFORMANCE	All data is stored in a deduplicated format.	It is highly compute intensive to rehydrate data from a pool of deduplicated data. Data Domain can take hours to days to do a restore, VM boot, or offsite tape copy.	95% to 98% of the total data restored comes from the most recent backup. Undeduplicated data is written to a disk landing zone for fastest ingest performance. The result is the most recent copy of the data is in the native backup format in a non-deduplicated format. Restores, VM boots, and offsite tape copies are 10 to 20 times faster than Data Domain. VM boots complete in single-digit minutes versus hours.

FUNCTIONALITY	DELL EMC DATA DOMAIN APPROACH	DELL EMC DATA DOMAIN CHALLENGES	EXAGRID ARCHITECTURAL APPROACH / BENEFITS
FIXED-LENGTH BACKUP WINDOW	Scale-up storage architecture with a fixed compute/processor/network ports front-end controller. As data grows, only disk capacity is added.	As the backup data grows, since the controller resources are fixed, it takes longer and longer to complete the deduplication. As data grows, the backup window continues to grow in length.	Full appliances, each with processor, memory, network ports and disk, are added as data grows - all resources are added into a scale-out system. The result is a backup window that stays fixed in length as data grows and protection of the initial investment, saving money and disruption over time.
DATA INTEGRITY	Does data integrity checking to ensure all the deduplicated data is present and not damaged or corrupted.		Does data integrity checking to ensure that all data is present and not damaged or corrupted.
REPLICATION TO DISASTER RECOVERY SITE, REPLICATION OVERHEAD	Replication uses compute during the backup window alongside inline deduplication.	Replication takes compute from the inline deduplication process, which slows down backups and expands the backup window. The controller has fixed compute, so when turning on a second process, the compute is reduced for inline deduplication.	Data is written to disk for fast backups. Deduplication and replication occur after data is committed to disk to not impede the backup window. Data is adaptively deduplicated in parallel with the backup window using multiple appliances in a scale-out system. Replication to the DR site does not impact ingest performance.
MULTIPLE DATA CENTER CROSS-REPLICATING SUPPORT	Supports replicating across major data centers so that each data center has an alternate DR data center target.		Supports replicating across major data centers so that each data center has an alternate DR data center target.
COST OF REPLICATION TO DISASTER RECOVERY SITE	Onsite system and offsite system must be the same size/capacity.	The cost of a two-site system is two times that of a primary site system.	Asymmetrical versus symmetrical: the second site system can be configured to be 100% repository; it is, therefore, half the capacity and cost of the primary site system, resulting in total cost for a two-site ExaGrid system at 1.5 times the primary site system versus 2 times - ExaGrid is lower cost.
MULTIPLE CLOUD DR OPTIONS	Can replicate to private, hybrid, and public clouds.		Can replicate to private, hybrid, and public clouds.
REPLICATION ENCRYPTION	Encrypts data sent over the WAN to the DR site.		Encrypts data sent over the WAN to the DR site.
REPLICATION BANDWIDTH	Moves only the changes, or 2% of the data, over the WAN to the DR site. Allows a bandwidth usage limit to be set for backup data replication.		Moves only the changes (or 2%) of the data over the WAN to the DR site. Allows a bandwidth usage limit to be set for backup data replication.
PERFORMANCE OF ENCRYPTED DATA AT REST	Encryption at rest is executed in software using the same compute resources as inline deduplication and replication.	Further slows down ingest on top of replication being turned on. Encryption takes compute away from ingest.	Encryption at rest is accomplished at the drive level. It does not take compute from ingest. The impact is negligible.
LOST TIME IN PRODUCTION DUE TO SYSTEM CLEAN UP	Requires periodic clean-ups and needs to be taken offline to run.	Takes the system out of production to do complete backups.	Does not include system clean-ups that take systems out of production.
ADVANCED INTEGRATIONS	Supports: • NetBackup / Backup Exec OST • Veeam SOBR • Oracle RMAN Channels	Parallel backups for performance is not supported due to a single front-end controller.	Support parallel backups into individual appliances in a scale-out system for increased performance and performance load balancing. In addition, supports Veeam Data Mover for increased ingest and fast synthetic fulls. Supports: • NetBackup / Backup Exec OST • Integrated ExaGrid-Veeam Accelerated Data Mover • Veeam SOBR • Oracle RMAN Channels
INSTALLATION AND TRAINING COSTS	Charges separately for installation and training.	Increases overall costs.	System is designed for easy plug-and-play installation with phone/WebEx support, eliminating installation charges. System is easy to use, and training occurs via WebEx in less than an hour at no additional charge.
CUSTOMER SUPPORT	Call into a pool of support techs	Need to work to get to a senior tech, need to repeat yourself on each call as each time you work with a different person. May or may not end up with a support that truly knows your backup application.	Assign a level 2 tech that is an expert on your backup application and knows your environment. Talk to the same tech every time. No need to fight to get to a senior tech or repeat yourself on every call. Monitoring of alerts and alarms at no additional charge for proactive support.

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