



ExaGrid Tiered Backup Storage

Fastest Backups.

Fastest Recoveries.

Unparalleled,
Cost-effective
Scale-out.

ExaGrid and Zerto Continuous Journal-Based Data Protection with Long-Term Backup Storage

To complete a business continuity and disaster recovery (BC/DR) plan it is critical that data is protected and recoverable during a disaster from very granular restore points. In addition, with the growing number of data protection regulations, such as HIPAA, GLBA, Sarbanes-Oxley, as well as the need to prepare for SEC audits and legal discovery, organizations must ensure that their IT environment is in compliance. This requires organizations to keep data for a long period of time such as months and years which means that the long-term retention of data must be protected both onsite and offsite. Zerto and ExaGrid have been working together to produce the best of both worlds, a solution that provides for granular real-time disaster recovery, as well as cost-efficient storage for long-term backup retention.

Zerto's Continuous Data Protection (CDP) captures and tracks modifications, automatically saving every version of the data the user creates locally at a target repository. This approach, in combination with ExaGrid's Landing Zone, Adaptive Deduplication process, and scale-out architecture, provides a solution that delivers continuous data protection and scalable long-term backup retention.

How Zerto's architecture integrates with the ExaGrid architecture:

- For short-term, real time data protection and disaster recovery, Zerto keeps a journal of all changes, providing granular recovery to just seconds before any issue.
- Zerto's Long-Term Retention capabilities with its Elastic Journal leveraging storage with data deduplication reduces the amount of storage required and controls the cost of long-term retention. ExaGrid offers the most efficient data deduplication achieving an average 20:1 duplication ratio.
- Using an ExaGrid system as a storage target for Zerto's Long-Term retention capabilities eliminates the issues associated with inline deduplication. Why? ExaGrid's unique architecture contains a front-end disk landing zone and a separate deduplicated data repository. Traditional deduplication appliances deduplicate the data inline resulting in slower backups due to the compute intensive process of data deduplication. Restores/recoveries are also slow since the data is only stored in a deduplicated format that needs to be rehydrated for each restore request.
- Zerto enables users to utilize the elastic journal to store data from any point in time for days, weeks, months, or even years which then can be sent directly to ExaGrid's landing zone for fast restores of recently added backups. In parallel with the backup, but not inline, ExaGrid deduplicates the data into the ExaGrid repository for long-term retention needs.
- Zerto's Data Protection Workflows automate recovery from any point in time, ensuring application consistency across VMs for both short and long-term data retention, in the ExaGrid system.
- As data grows, appliances are simply added to the ExaGrid system which eliminates controller forklift upgrades and product obsolescence, versus traditional scale-up solutions that force forklift upgrades and force periodic product obsolescence.
- ExaGrid can replicate all data in the onsite ExaGrid system to an offsite location, public cloud, or a secondary physical ExaGrid system, ensuring that all long-term retention data is protected from a site disaster.

Together, ExaGrid and Zerto provide the best of both worlds:

- Continuous data protection with fast restores for Disaster Recovery and Backup
- Cost-effective long-term retention storage
- Intelligent Index and Search of all protected data
- Low-cost, offsite long-term retention storage

Scenarios Supported by the ExaGrid and Zerto Solution:

