

ExaGrid Disk Backup with Deduplication Competitive Comparison

Category	ExaGrid	Competitors
Deduplication Ratio	<ul style="list-style-type: none"> ✓ Ratios ranging from 10:1 to 50:1 <ul style="list-style-type: none"> • Byte-level delta: only bytes changed from previous backup are stored 	<ul style="list-style-type: none"> ✓ Ratios ranging from 10:1 to 50:1 <ul style="list-style-type: none"> • Block-level : breaks backup jobs into 8KB blocks and compares hashes to find redundant blocks.
WAN-efficient DR	<ul style="list-style-type: none"> ✓ Only changed backup data replicated offsite <ul style="list-style-type: none"> • “Instant DR “ via rebuild from most recent backup 	<ul style="list-style-type: none"> ✓ Only changed backup data replicated offsite <ul style="list-style-type: none"> • Slower recovery from “rehydrating” deduped backup
Virtualized Environment Support	<ul style="list-style-type: none"> ✓ Achieves faster backups and restores for VMs <ul style="list-style-type: none"> • Fast backups of VMs via post-process • Rapid recovery, with most recent VMs intact 	<ul style="list-style-type: none"> ✗ Works in Virtual Environments <ul style="list-style-type: none"> • In-line deduplication takes longer to protect VMs • Slower VM recovery, as images need to be rebuilt
Shortest Backup Window	<ul style="list-style-type: none"> ✓ Fastest, via post-process deduplication <ul style="list-style-type: none"> • Backups at disk speed • No bottleneck in fast environments 	<ul style="list-style-type: none"> ✗ Slower, via in-line deduplication <ul style="list-style-type: none"> • Deduplication throttles backups • Bottleneck in fast environments
Fastest Restores and Tape Copies	<ul style="list-style-type: none"> ✓ Most recent copy intact <ul style="list-style-type: none"> • Restores at disk speeds • Fast tape copies 	<ul style="list-style-type: none"> ✗ Most recent copy is the most fragmented <ul style="list-style-type: none"> • Millions/billions blocks reassembled per TB • Slower restores/tape copies
Backup storage capacity	<ul style="list-style-type: none"> ✓ Largest singly managed system: 100 TB <ul style="list-style-type: none"> • Maximum throughput 18 TB/hour (100TB GRID) 	<ul style="list-style-type: none"> ✗ Largest system: 60TB <ul style="list-style-type: none"> • Maximum throughput 5.4TB/hour (60TB system)
Scalability: Maintain Short Backup Times	<ul style="list-style-type: none"> ✓ Strong: add appliances (disk, CPU, mem, GigE) <ul style="list-style-type: none"> • Backups/restores remain fast • Automatic capacity load balance • Data deduplication shared across the GRID 	<ul style="list-style-type: none"> ✗ Weak—can only add disk <ul style="list-style-type: none"> • Low end – rip and replace to maintain backup times • High end – backups/restores take longer • Data deduplication not shared across isolated systems
Scalability: No forklift upgrades	<ul style="list-style-type: none"> ✓ Grow via GRID in cost-effective increments <ul style="list-style-type: none"> • Never rip and replace to add capacity • Preserves investment 	<ul style="list-style-type: none"> ✗ Forklift upgrades required to expand <ul style="list-style-type: none"> • Outgrow max system size, must buy new controller • Controller is 70 to 80% of cost
Management and Reporting	<ul style="list-style-type: none"> ✓ Robust and Content Aware <ul style="list-style-type: none"> • Manage multiple sites and devices via single UI • Dedupe ratio and replication status by job 	<ul style="list-style-type: none"> ✗ Complex management and generic reports <ul style="list-style-type: none"> • Configuration via command-level (CLI) per device • Dedupe ratios and replication status by device
Cost	<ul style="list-style-type: none"> ✓ Best price for highest performance <ul style="list-style-type: none"> • Cost-effective GRID architecture 	<ul style="list-style-type: none"> ✗ Higher price for lower performance <ul style="list-style-type: none"> • CPU-intensive, requires expensive hardware