



Kashya KBX5000 Lowers Costs of Advanced Data Protection and Recovery

Analyst: Michael Fisch

Management Summary

The Kashya *KBX5000 Data Protection Appliance* is a solution for enterprise data protection with multiple capabilities. It is more challenging to describe because it does not fit easily into one product category. Rather, it spans several.

A smart way to evaluate a solution like this is first to have an understanding of your enterprise's recovery requirements. How quickly and fully do you need to recover data after a system failure or disruptive event? Do you need to replicate to a remote site to recover from a local disaster? The answers help determine which data protection technologies your enterprise should deploy – and thus whether the KBX5000 may be a good fit.

The KBX5000 is a SAN-resident appliance available in software modules for Remote Replication and Continuous Data Protection (CDP). It offers multiple data replication and copying technologies:

- **Remote mirroring** – Maintains a full, real-time copy of data at a target location over an unlimited distance (synchronous or asynchronous).
- **Snapshot (point-in-time) copy** – Creates data copies at specific points in time.
- **Continuous data protection** – Recovers data to virtually any prior point in time.

In addition, it has a number of supporting features that increase performance and/or lower costs:

- **Data reduction** – Reduces data as much as 25:1 (typically 8:1) using multiple techniques, lowering communication and storage costs.
- **Heterogeneous support** – Protects data across multiple, diverse servers and storage arrays.
- **Consistency groups** – Establishes application-related volume groups that simplify management.
- **Application-aware bookmarks** – Correlates application events (e.g., checkpoints, cache flush, hot backups) to specific points in time within the CDP journal, so users can more easily pick optimal points for fast restarts.
- **Fibre Channel to IP routing** – Displaces the need for edge routers.
- **Adaptive replication** – Dynamically changes replication modes to meet user-defined policies for lag time or bandwidth consumption.

In short, the Kashya KBX5000 does a lot for enterprise data protection. It is especially suitable for enterprises with SANs, heterogeneous servers and storage, and more stringent recoverability requirements. Read on for the details.

IN THIS ISSUE

➤ Business Continuity	2
➤ Kashya KBX5000	2
➤ Business Benefits	3
➤ Conclusion	3

Business Continuity

Information has always been essential to doing business, and never more so than today. When enterprise data is lost or unavailable due to a system failure or disruptive event, workers cannot perform basic tasks, like process an order, receive a customer support call, work on an engineering project, or send e-mail. Business activities slow or cease. Therefore, **business continuity depends on maintaining data access, and data access depends on effective protection and recovery**¹.

This is why enterprises are increasingly willing to invest in data protection technologies that provide faster and more complete recovery. It is an investment in the productivity and continuity of the enterprise.

What level of recoverability does a given enterprise need? It depends on business requirements – operations, corporate governance, and industry regulations. Moreover, not all data is equal in terms of relative value and need for protection. Some is critical; some is moderately important; some only needs to be archived “just in case”. Enterprises should analyze and understand their particular requirements. Then they may consider which technologies would best serve their interests.

Kashya KBX5000

The Kashya *KBX5000 Data Protection Appliance* is an advanced solution for enterprise data protection. The appliance consists of specialized software running on industry-standard server hardware, which can be configured in an active-active cluster for high availability. The appliance connects directly to a Fibre Channel SAN. It receives mirrored writes from client servers² as the source for copying and replication, which it stores locally and/or sends over an IP WAN to a target appliance at a remote site. The KBX5000 is *out of band*, meaning that it does not sit in the data path, nor create a potential point of failure. Enterprises that are sensitive to high availability may prefer this approach. Kashya offers two modules for the KBX5000, both of which can run on the same appliance: the

¹ See **The Clipper Group Explorer** dated January 28, 2005, entitled “*Recovery in Perspective – Ensuring Access to Enterprise Data*” and available at <http://www.clipper.com/research/TCG2005003.pdf>.

² Client servers run a filter driver for mirroring writes. Alternatively, it can use the *SANTap* protocol in the *Cisco MDS 9000* director or switch to avoid affecting clients.

KBX5000 CDP for local data protection and the *KBX5000 CRR* for remote replication.

The KBX5000 offers all of these capabilities for data replication and copying.

- **Remote mirroring** – Maintains a full, real-time copy of data at a target location over an unlimited distance. The operation can be synchronous (no lag between the primary and target) or asynchronous (lag of seconds to minutes). It provides a form of synchronous replication over unlimited distances by buffering the most recent writes at the local or another metro-area site before transmitting them over the WAN.³ The mirroring capability is bidirectional.
- **Snapshot (point-in-time) copy** – Creates copies of data at specific points in time, which may be stored locally and/or replicated to the remote site. If mirroring is analogous to a window with a real-time view, a snapshot copy is like a photograph capturing a scene at a fixed point in time. Snapshots are useful for non-disruptive backups, quick recovery to a previous point, and application testing and development. The KBX5000 provides space-saving differential copies. It also has a special technology called small-aperture snapshot that maintains copies only seconds apart.
- **Continuous data protection (CDP)** – Recovers data to virtually any prior point in time. CDP is a more recent technology than mirroring and snapshots. If mirroring is analogous to a window and snapshots a photo, then CDP is like video. The KBX5000 stores a continuous log of changes and can create a read/write copy at any prior point in time. Enterprises can recover more quickly logical errors, like data corruption or viruses, by rebuilding from a point just prior to the disruption. The CDP journal may be stored locally or replicated to a remote site.

In addition, the KBX5000 provides a number of supporting features that enhance the performance or lower the costs.

- **Data reduction** – Reduces data as much as 25:1, though more typically in the range of 8:1, by employing several complementary techniques: *delta differentials* that replicate changed data at a more granular level than

³ Normally, synchronous mirroring is limited to distances of about 100 km because transmission latency slows application performance.

whole blocks, *hot spot reduction* that replicates only most recent data changes, proprietary *data compression*, and *application-specific buffering* for *Oracle*, *SQL server*, and *DB2*. Data reduction increases replication speeds and lowers bandwidth and storage requirements.

- **Heterogeneous support** – Protects data across multiple, heterogeneous servers and storage arrays. Many competitive replication and copying technologies run on specific servers or storage arrays, which limit their scope, whereas the KBX5000 can provide a single data protection solution for all SAN-connected servers and storage.
- **Consistency groups** – Creates logical, application-related volume groups. Consistency groups allow all volumes related to an application to be handled as a single entity, even if the application spans multiple hosts or the volumes multiple storage arrays. Enterprises can use them to set protection and recovery policies for specific applications, based on cost-performance tradeoffs and business requirements.
- **Application-aware bookmarks** – Establishes I/O-based tags that correlate application events to specific points in time within the CDP history journal. They allow enterprises to create a series of time- or event-driven consistency points when used with application-specific features, like Oracle hot backup or Microsoft *Volume Shadow Copy Services* (VSS). The result is faster and easier recovery.
- **Fibre Channel to IP routing** – Routes data between a Fibre Channel SAN and IP WAN, displacing the need for edge routers.
- **Adaptive replication** – Changes replication modes dynamically and automatically to comply with user-defined policies for lag time or bandwidth consumption. Depending on business requirements, the policies can reflect the need for a particular degree of recovery (in terms of RPO) or a limitation on bandwidth to avoid affecting the performance of other applications sharing the WAN link. The KBX5000 is unique in its ability to switch between synchronous and asynchronous mirroring and snapshots.

KBX5000 pricing is based on the amount of protected storage. Entry configurations for the KBX5000 CDP start at \$50,000 and the KBX5000 CRR starts at \$65,000.

Business Benefits

The feature set of the Kashya KBX5000 offers several business benefits.

Lower Hardware and Software Costs

It can lower hardware and software costs for data protection in several ways. First, the KBX5000's multiple copying and replication technologies and supporting features allow it to displace a "basket" of alternative products, such as mirroring software, snapshot software, CDP, edge routers, and management software. Since alternative products also may be limited to specific storage arrays or operating systems, its heterogeneous support expands the potential for displacement. Furthermore, enterprises are not locked in to particular storage arrays, so, for instance, they could use a low-cost SATA array as a replication target. The data reduction feature lowers the amount of capacity required for replication and, therefore, storage costs.

Lower Operating Costs

A single data protection solution is simpler to manage than multiple, disparate solutions, which reduces ongoing management costs in the form of IT staff-hours. Data reduction lets one use a lower-bandwidth WAN link, reducing a big component of replication costs. Adaptive replication helps get the best utilization from a WAN link while meeting business requirements.

Enhanced Recovery

Data reduction, consistency groups, application-aware bookmarks, adaptive replication, and CDP affect recovery in different ways, although all are positive. They can improve both recovery point objective (RPO) and recovery time objective (RTO). Enhanced recovery, once again, translates into business continuity and productivity.

Conclusion

The Kashya KBX5000 targets enterprises with SANs that need more advanced – and perhaps multiple – data protection and recovery capabilities. It fits well in distributed, heterogeneous environments. **Be sure to consider it if your environment is like this, and especially if you are looking for ways to lower data protection costs.**



About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Michael Fisch is Director of Storage and Networking for The Clipper Group. He brings over ten years of experience in the computer industry working in sales, market analysis and positioning, and engineering. Mr. Fisch worked at EMC Corporation as a marketing program manager focused on service providers and as a competitive market analyst. Before that, he worked in international channel development, manufacturing, and technical support at Extended Systems, Inc. Mr. Fisch earned an MBA from Babson College and a Bachelor's degree in electrical engineering from the University of Idaho.

- ***Reach Michael Fisch via e-mail at mike.fisch@clipper.com or at 781-235-0085 Ext. 211. (Please dial "211" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, and "***clipper.com***" are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "***Navigating Information Technology Horizons***", and "***teraproductivity***" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.