Advances in Data Protection and Replication

A Kashya Technical White Paper

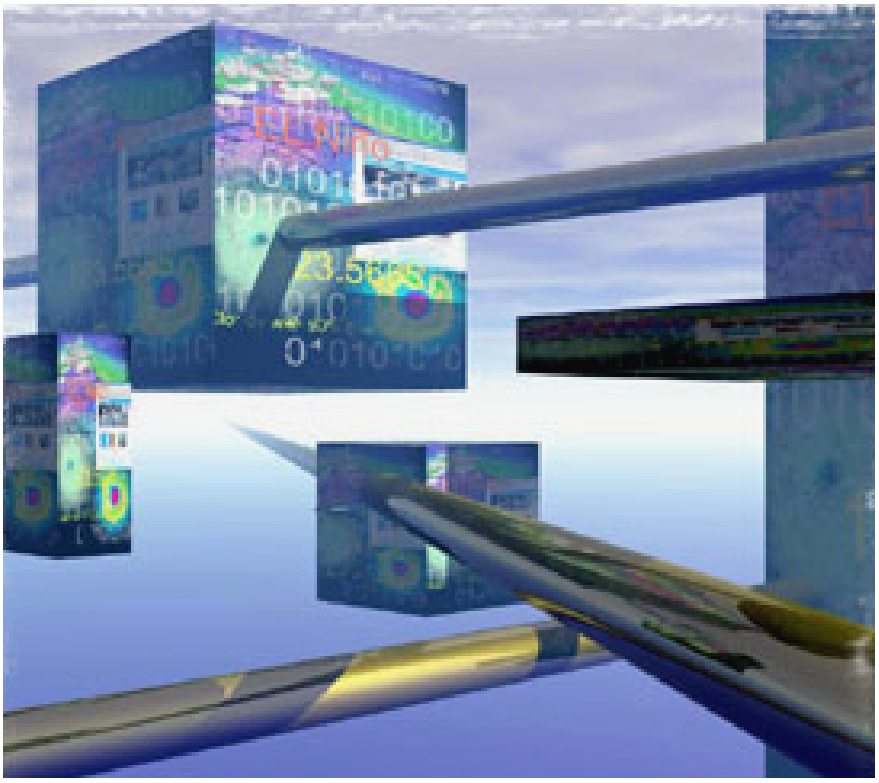

Kashya™ develops unique algorithmic technologies to enable an order of magnitude improvement in the reliability, cost, and performance of an enterprise's data protection capabilities. Based on the Kashya Data Protection Appliance™ platform, Kashya's powerful solutions deliver superior data protection at a fraction of the cost of existing solutions. Kashya's Data Protection Appliance connects to the SAN and IP infrastructure and provides bi-directional replication across any distance for heterogeneous storage, SAN, and server environments.

I. The Challenge

The data that drives today's globally oriented businesses is stored on large networks of interconnected computers and data storage devices. This data must be 100% available and *always* accessible and up-to-date, even in the face of local or regional disasters. Moreover, these conditions must be met at a cost that is affordable, and without in any way hampering normal company operations.

An up-to-date replica of a company's mission-critical data at a secondary site is a prerequisite for uninterrupted business operations following a disaster. Any local condition that damages data at the primary site must not be able to harm the replica of that data.

Distance is the key to reducing the risk. Power outages, fires, floods, earthquakes, wars, and, yes, malicious attacks, all can temporarily or permanently halt normal business operations where they strike. The industry watchdogs such as SEC and HIPAA have published guidelines that recommend a minimum of 300 miles separation between the primary and secondary site.

To reduce the business risk of an unplanned event of this type, an enterprise must ensure that a copy of its business-critical data is stored at a *secondary* location. Synchronous replication, used so effectively to create perfect copies in *local* networks, performs poorly over longer distances (see Figure 1, "Replication Methods").

Figure 1: Replication Methods

Synchronous – *Every write transaction committed must be acknowledged from the secondary site*. This method enables efficient replication of data within the local SAN environment. Extending this approach to long-distance copy over the WAN, however, is both very costly and dramatically degrades the performance of critical applications.

Asynchronous – *Every write transaction is acknowledged locally and then added to a queue of writes waiting to be sent to the secondary site.* With this method, some data will normally be lost in the event of a disaster. This requires the same bandwidth as a synchronous solution. While it is less disruptive to many applications, performance will deteriorate to that of synchronous replication for "write-intensive" applications, or in the event of even the shortest WAN outage. Finally, in most cases, asynchronous replication introduces the problem of data inconsistency. Most current solutions in the market can not guarantee "write-order fidelity" at the secondary site, and even the most advanced solutions cannot retain write order in a multi-host, multi-storage environments. Further, if a disaster occurs during resynchronization (i.e., following a temporary malfunction), no consistent copy will exist.

Snapshot – A *consistent image of the storage subsystem is periodically transferred to the secondary site. Only the changes made since the previous snapshot must be transferred, resulting in significant savings in bandwidth.* By definition, this solution produces a copy that is not up-to-date; however, increasing the frequency of the snapshots can reduce the extent of this lag. Many of today's snapshot implementations, however, are prohibitively expensive, as they transfer data at large granularity, and demand as many as five times the production site storage in order to guarantee a single copy of the data at the secondary site. Snapshot is also known as "point-in-time (PIT)" or "checkpoint" replication.

Small-Aperture Snapshot – Kashya's system offers the unique ability to take frequent snapshots, just seconds apart. This innovative feature is utilized to minimize the risk of data loss due to data corruption that typically follows rolling disasters.

A copy of critical company data at a secondary site is also a prerequisite for continuity in business operations. For an organization that already has multiple geographically dispersed installations, one of these locations is a logical first choice for the secondary site. A disaster recovery solution must be able to "bridge" the distance. The system must support *rapid* failover to the secondary site, and it must support transparent failback following correction of the problem at the primary site.

II. Requirements

Requirement 1:
The data protection solution must provide maximum flexibility in the definition and implementation of company business priorities.

Companies have different methods for expressing policies on issues such as resource allocation and performance. One company may focus more on bandwidth consumption than on the *lag* between the local and secondary copies. A second may be primarily concerned with the projected cost of potentially losing data in the event of a disaster. A third may use yet some other set of parameters, as most appropriate for its particular business requirements. In addition to offering multiple data protection policies to support different and dynamic business priorities, the solution must include the ability to automatically adjust the behavior of the system based on real-time application workload and bandwidth availability, to maximize protection with the available resources while adhering to the user-defined policies.

Requirement 2:
The copy must be up-to-date, enabling failover with no data loss.

In LAN and SAN environments, synchronous replication is used to create perfect copies with little or no overhead. At longer distances, however, neither synchronous nor any other replication method alone can provide the perfect solution for up-to-date copies. The solution must be able to produce up-to-date copies at the required distances in support of an effective business continuity strategy.

Requirement 3:
There must be at least one reliable copy of the data at all times.

Inevitably, often due to human error, data becomes *corrupted*. The corruption will be detected sooner or later, through human intervention or perhaps by application-specific tools developed for this purpose. Once detected, however, the system should be able to eliminate the corruption by restoring the most up-to-date uncorrupted copy of the data. For this purpose, the replication solution must provide a mechanism by which the uncorrupted version of the secondary copy can be readily recovered.

Requirement 4:
Replication must not cause degradation in normal business operations.

Business-critical applications must not be halted in order to create a copy. Moreover, the solution must not steal resources from the host and storage systems therefore impacting applications. Replication must occur without hampering the performance of the applications. The system must offer synchronous levels of protection with no application degradation, even over long distances.

Requirement 5:
Use of network resources for data replication must be optimized.

The system must satisfy data replication and disaster recovery requirements while minimizing the amount of data transferred. In snapshot replication, gathering data at the local site can yield bandwidth savings that are not possible with synchronous and asynchronous replication. In addition to reducing bandwidth requirements, the system must make maximum use of *existing* communications equipment and infrastructure, including use of the existing IP network infrastructure.

Requirement 6:
Use of additional storage resources for data replication shall be minimized.

The system must avoid the requirement for large quantities of additional storage related to the replication process. Solutions that entail buffering, or intermediate storage of data between the local and secondary sites should store *only* the modified data (rather than copies of entire volumes), to the extent possible.

Requirement 7:
All transfer of data must be secure.

Critical data is often a company's most valuable asset. Consequently, the system must secure company information assets against all unauthorized access during and after transfer of data from the primary to the secondary site.

Requirement 8:
The solution must readily facilitate testing of DR capabilities at the secondary site.

A disaster recovery plan must be tested on a regular basis, to ensure that all aspects of the plan have been carefully addressed. The system must enable direct Read/Write access to data at the secondary site, without the need to first make an additional copy and without disrupting the replication process.

Requirement 9:
The solution must be universal, scalable, and always available.

> The system must have no operational dependencies on any particular application, server platform, networking infrastructure, and storage platform; it should support any combination of existing and future software, hardware, and operating systems. In this way, companies can leverage their existing IT infrastructure investment (from any combination of vendors), and reduce costs. The system must also offer the ability to grow incrementally, and only when needed, to respond to changing needs. It must also be always available to support non-stop business operations.

Current data protection solutions have been unable to adequately satisfy all of these requirements — some of which are seemingly contradictory — especially over longer distances. Those that provide true up-to-date replication (through synchronous replication) cannot support the longer distances without severely degrading the performance of the host applications. Those that most effectively conserve bandwidth are not absolutely up-to-date. None allows the system to change the replication method dynamically in response to real-time application and network activity.

III. Defining the Solution

1. The Kashya Approach

Kashya delivers data protection and business continuity solutions that are:

- **Universal** – supporting heterogeneous server and storage platforms, enabling a complete data protection solution for the entire enterprise.
- **Reliable** – delivering real-time, always consistent replicas of critical data, over any distance to guarantee immediate, no-data-loss recovery of company operations.
- **Bandwidth-Efficient** – employing application-aware and storage-aware algorithmic bandwidth and data reduction technologies that deliver unprecedented efficiencies in bandwidth use.
- **Storage-Efficient** – utilizing unique snapshot and buffering technologies to eliminate the need for intermediary and additional copies while allowing replication to any storage allowing lower deployment of the lower cost storage at the secondary site.
- **Flexible** – providing advanced policy-based capabilities to optimally and autonomously balance cost with functionality and performance, in a way that meets the unique business needs of each application in the enterprise.
- **Up-to-Date Protection** – delivering up-to-date, synchronous levels of protection with no distance limitation and without application degradation.

- **Manageable** – providing a single-point of management for any number of data protection appliances through an intuitive Graphical User Interface (GUI) and scriptable Command Line Interface (CLI).
- **Available** – enabling uninterrupted availability of data protection services.
- **Scalable** – supporting customer needs by linearly scaling performance in line with any increase in data rates and size of replicated storage.

2. *Kashya Architecture*

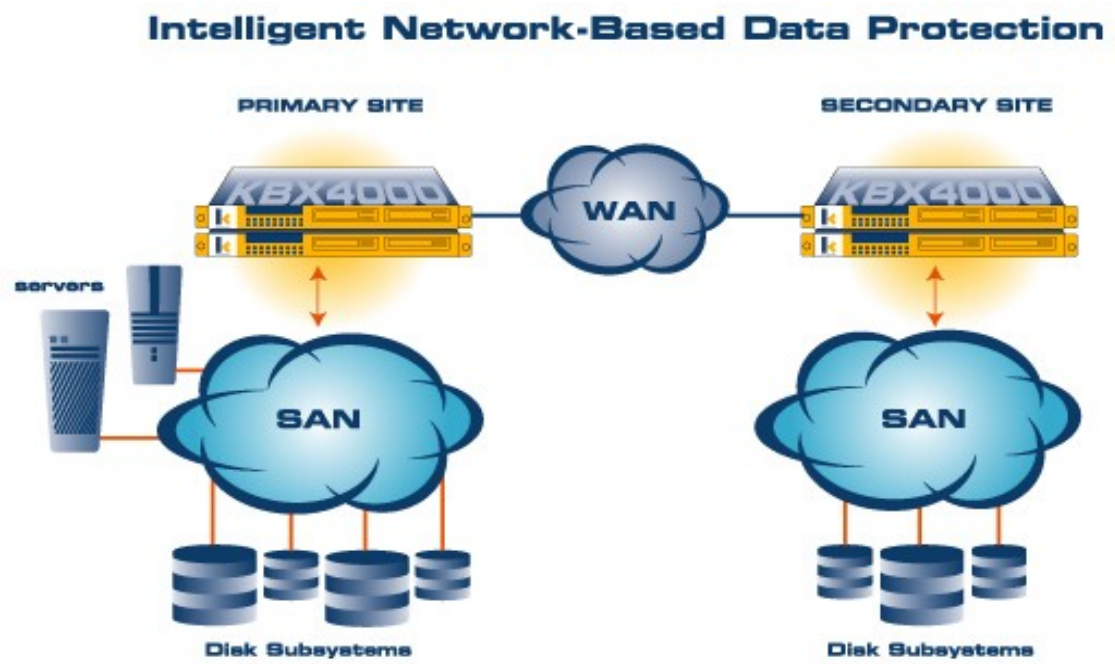

Figure 2: System Architecture

At the core of Kashya's architecture are a number of *advanced patent-pending technologies* that, when deployed with Kashya's *innovative architecture and design*, unleash powerful data protection solutions at a fraction of the cost of existing products.

Kashya's advanced architecture can be summarized as follows:

- *Positioning at the junction between the SAN and the IP infrastructure* enables Kashya solutions to:
  - Deploy enterprise-class data protection non-disruptively and non-invasively
  - Support heterogeneous server, SAN, and storage platforms
  - Monitor SAN and WAN behavior on an ongoing basis, to maximize the data protection process.
- *Advanced algorithms*, that:
  - Automatically manage the replication process, with strict adherence to user-defined policies that are tied to user-specified business objectives

- Use real-time information to enable the system to dynamically adapt its replication mode based on SAN and WAN conditions, application activity, and on user-defined policies
- Dramatically reduce the requirements for both storage and WAN bandwidth.

- *Cluster of independent, intelligent Kashya appliances, with unique data handling capability*, enabling the Kashya solutions to:
  - Deliver synchronous level protection with no intermediary copies of the data, with no distance limitations, and without the associated degradation of host application performance
  - Guarantee a consistent replica of business-critical data in the event of any failure or disaster
  - Enable incremental, non-disruptive growth and always-on availability.

IV. Delivering the Solution

The *Kashya Data Protection Appliance*™ provides superior data protection — in the face of any disaster, striking at the local site, the secondary site, or anywhere in between — at a fraction of the total cost of other solutions.

1. Dynamic Support for Business Policies

Requirement 1: The data protection solution must provide maximum flexibility in the definition and implementation of company business priorities.

The Kashya Data Protection solution offers unparalleled flexibility in implementing the widest possible range of business policies. The Kashya system offers innovative policies for balancing data protection with cost, based on each organization's unique set of business and technical criteria. These user policies are defined at the level of a *consistency group* (a consistency group is a logical grouping of hosts and volumes associated with one or more applications), allowing organizations to set specific policies for different applications. Applications may be in one or more hosts, and on one or more volumes from one or more storage devices on the SAN. Furthermore, volumes may be housed in one or many heterogeneous SAN storage devices. Kashya *consistency groups* thus provide a high degree of flexibility, allowing the Kashya system to react to ever-changing conditions with regard to the data storage/transfer load and bandwidth availability, while maintaining the user-specified policies automatically without any user intervention.

Other products require that the user specify the replication method; i.e., synchronous, asynchronous, or snapshot (see Figure 1, *Replication Methods*) at the time of configuration. For snapshots, it is also necessary to specify the snapshot frequency that is to be used for all data replication. Among the products that offer snapshot replication, none offers either of the other replication methods. Such an inflexible approach is incompatible with real-life situations, in which application activity and bandwidth availability can change dramatically over time.

In contrast, the Kashya Data Protection Appliance offers *a full spectrum of replication modes*, from synchronous, to asynchronous, to snapshot, to small-aperture snapshot. The replication process is managed automatically, with strict adherence to user-defined policies that are tied to desired business objectives. *The system adapts its replication modes dynamically*, based on the available bandwidth and the application workload, to achieve the stated business objectives for each application. This greatly simplifies data and disaster recovery management for complex and heterogeneous environments.

For example, if the user's primary objective is to *minimize lag*, the system will use as much bandwidth as is available, in an attempt to keep the time offset as small as possible between source and copy. In all cases, the system ensures that the *maximum permissible* lag between sites is not exceeded.

Alternatively, the user can select a *minimize bandwidth* policy for less critical data causing the system to use as little bandwidth as possible, while maintaining the lag within the specified lag limit.

2. Efficient Utilization of Network Resources

❑ Intelligent use of bandwidth

Requirement 5:
Use of network resources for data replication must be optimized.

The Kashya Data Protection Appliance employs intelligent bandwidth reduction technologies that deliver unprecedented reduction in bandwidth requirements, dramatically reducing WAN costs, particularly over long distances. Data reduction is achieved through application-aware and storage-aware algorithmic techniques that conserve bandwidth to the extent not possible with traditional compression technologies.

■ Hot spot data reduction

The K-Box identifies any block of data that has been modified more than once during a snapshot period, and sends only the most recently revised data over the WAN. When write transactions overlap, Kashya subdivides the modified data to ensure that it sends only the most-recent changes to the data (i.e., including where the write transactions overlapped).

■ Delta differentials

Many applications write data to disk in blocks as big as 8K or 16K, even though only a few bytes in the block have actually have been modified. Oracle, for example, can be configured to write in blocks anywhere from 2K to 64K, with 8K a popular choice for many applications. Kashya *delta differential* technology achieves significant savings in bandwidth by examining the modified data at finer granularity, and sending only the data that has actually been modified.

■ Algorithmic compression

Kashya has developed a proprietary data compression algorithm that achieves superior compression with reduced CPU requirements.

- **Application-specific buffering**

For select applications, Kashya achieves even better results by tailoring proprietary compression techniques to the characteristic data found in each application. Applications based on Oracle, SQL server, and DB2 have yielded results that are at least twice as good as those achieved with the best compression technologies otherwise available on the market.

Examining the Cumulative Effects of Bandwidth Reduction Technologies

In Figure 4, the bar charts show the cumulative effects of the Kashya techniques used for bandwidth reduction on Oracle data and log files, as created by the TPC-C benchmark. Snapshots of 250 MB and 10 GB were tested. The extent of the bandwidth reduction achieved by these techniques varies according to the types of data being transferred and the size of the snapshots. "*Algorithmic Optimization*" refers to the Kashya proprietary compression algorithm; "*Application Optimization*" refers to compression tailored the particular application.

The chart shows bandwidth reduction rates ranging from 1.5× (33%) to 17.3× (94%). Note that the impact of hot spot analysis grows dramatically from 1.5× (33%) in the smaller file to 3.2× (69%) in the larger, reflecting the greater number of overwrites to the same disk location in the larger snapshot.

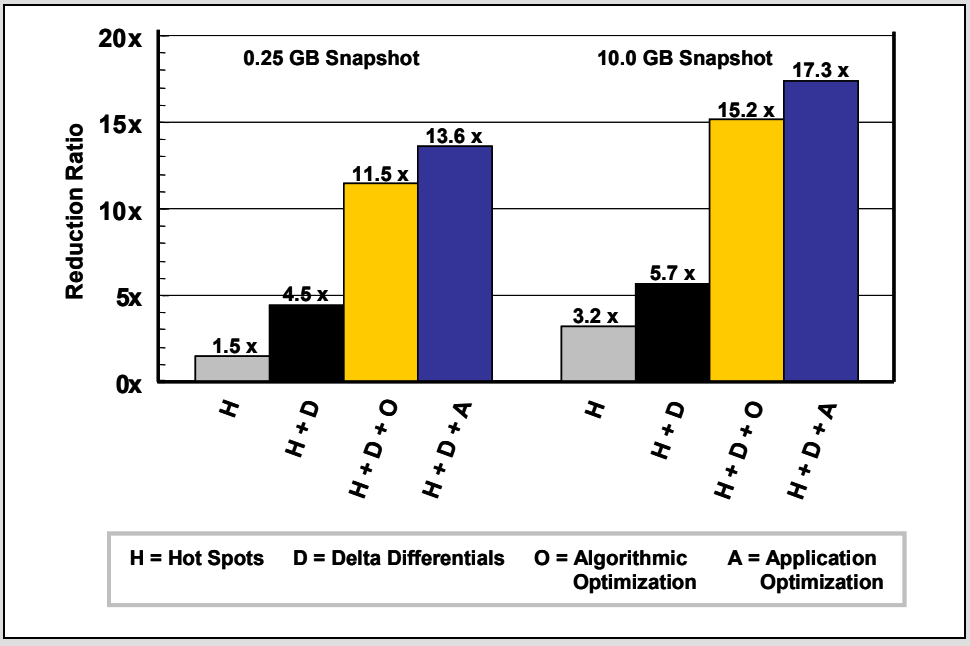

Figure 3: Bandwidth reduction, transmission of Oracle data and file system logs

Dynamic use of WAN network resources

The Kashya Data Protection Appliance uses the IP network in the most efficient manner for each type of replication policy. A proprietary algorithm balances the behavior of multiple internal connections to the available IP transport resources, on the assumption that these resources are constantly in flux, even over very short periods of time. Consequently, Kashya is able to use every type of WAN connection, from dedicated lines to unreliable public IP networks.

Other replication products use separate components for replication and fiber channel-to-IP conversion. The fact that Kashya integrates replication logic and fiber channel-to-IP conversion in the same component yields several significant advantages:

- When transmitting synchronously, the Kashya system can combine metadata and data in the same transmission over the WAN, thereby cutting by *half* or even *two-thirds* the normal latency created by synchronous replication through a standard fiber channel-to-IP converter.
- In low-lag situations (i.e., using synchronous replication), the connection can be optimized for low latency. In high-lag situations, the connection can be optimized with regard to bandwidth utilization.

3. Technologies for Disaster Recovery

Efficient buffering

Requirement 4:
Replication must not cause degradation in normal business operations.

Every time a host application writes to the local disk storage subsystem it writes it in parallel to the local Kashya appliance. Kashya uses this synchronous connection together with its exclusive buffer to deliver synchronous, up-to-date levels of protection with no application degradation to any distance, and without the need for additional storage.

Requirement 2:
The copy must be up-to-date, enabling failover with no data loss.

In the event of a primary disk failure, the Kashya system enables failover by flushing its buffer to the secondary site, with absolutely no loss of data. The Kashya buffer can be stored on the Kashya appliance, which can be located fibre-channel distances away, provided this copy of the data is not in close proximity to the local/primary data for better protection against local disasters.

With other DR solutions, failover with no data loss is possible only when using synchronous replication. Synchronous replication, however, is costly to implement and requires high bandwidth. In addition, as latency increases with distance, synchronous replication can negatively affect the performance of the application.

Requirement 6: Use of additional storage resources for data replication shall be minimized.

The Kashya Data Protection Appliance does not require large amounts of extra storage in order to hold multiple copies of the data during replication. Rather, the Kashya appliance holds *only the data that has been modified*. Kashya's bandwidth reduction technologies will further reduce storage requirements.

- **Multiple snapshot preservation**

Requirement 3: There must be at least one *reliable* copy of the data at all times.

The buffers in the Kashya appliance provide superior protection against data corruption that can occur in the event of rolling disasters. Kashya's data-handling efficiency enables it to save several consecutive snapshots at the secondary Kashya appliance without a large storage overhead. In the event of a rolling disaster, the operator at the secondary site checks back through the recent snapshots, eliminates corrupted snapshots, and initiates the reassembly of the secondary copy of the data from the most-recent uncorrupted snapshot. Moreover, because the Kashya system buffers only the modified data, this method can be achieved without requiring large amounts of additional storage.

- **Guaranteed data consistency without additional copies**

Requirement 3: There must be at least one *reliable* copy of the data at all times.

Data processed by the Kashya Data Protection Appliance remains consistent at the secondary site in the face any system failure, whether at the local Kashya appliance or storage subsystem, over the WAN network, or at the secondary Kashya appliance or storage subsystem. Moreover, Kashya accomplishes this without making additional copies of the data.

4. Heterogeneous Hosts and Storage Subsystems

Requirement 9: The solution must be universal, scalable, and always available

The Kashya Data Protection Appliance is uniquely situated to support a system configuration that includes heterogeneous host, SAN and storage platforms. In contrast, storage-based data protection solutions only support that individual storage hardware platform.

As data protection needs grow, Kashya appliances can be added non-disruptively, resulting in unlimited scalability. In addition, the Kashya system employs a cluster of multiple, active-active independent appliances that can be added, upgraded, or replaced without affecting the replication process, resulting in always-on availability.

5. Data Processing at the Secondary Site

Requirement 8:
The solution must readily facilitate testing of DR capabilities at the secondary site

The Kashya Data Protection Appliance enables direct Read/Write access to data at the secondary site, without the need to first make an additional copy and without disrupting the replication process. As a result, disaster recovery plans can be tested at the secondary site on a regular basis, without requiring additional storage, and without impacting on-going replication. Other workloads that can be deployed in the same non-disruptive manner at the secondary site include development and backups.

6. Technologies for Data Security

Requirement 7:
All transfer of data will be secure.

Kashya uses standard data security methods to safeguard the storage and transfer of the data with which it is entrusted. Users decide whether data is to be encrypted or not. If so, Kashya offers a choice between encryption of all replicated data, or of just the data headers, which supports higher throughput. Moreover, this parameter is configurable on a per-volume basis.

V. Kashya - Enterprise Data Protection…with No Limits

The *Kashya Data Protection A*ppliance guarantees an up-to-date copy of your company's data at a secondary site, and supports an immediate no-data-loss failover of company operations to that site — for complete data protection in the face of any disaster.

Functional Requirement	Kashya Complies?	Comments
Requirement 1: The data protection solution must provide maximum flexibility in the definition and implementation of company business priorities.	✓	Kashya Data Protection Appliance *offers policy-based replication and manages the replication process automatically, with strict adherence to user-defined policies that are tied to desired business objectives. The system adapts its replication dynamically, based on the available bandwidth and the application workload.*
Requirement 2: The copy must be up-to-date, enabling failover with no data loss.	✓	Kashya Data Protection Appliance *provides a suite of technologies that makes sure you get the most up-to-date copy possible for your disaster-recovery investment.*
Requirement 3: There must be at least one *reliable* copy of the data at all times.	✓	Kashya Data Protection Appliance *provides maximum protection against data corruption due to human errors and rolling disasters. Moreover, the system guarantees a consistent replica of business-critical data in the event of any failure or disaster.*

Functional Requirement	Kashya Complies?	Comments
Requirement 4: Replication must not cause degradation in normal business operations.	✓	*Kashya's architecture is based on a cluster of independent appliances. It does not reside on the host server, where it would use expensive memory and CPU cycles, and not on the storage subsystem, where it would waste expensive disk resources. The* Kashya Data Protection Appliance *delivers synchronous levels of protection with no application degradation and no distance limitations.*
Requirement 5: Use of network resources for data replication will be optimized.	✓	Kashya Data Protection Appliance *minimizes bandwidth requirements through innovative data reduction technologies and adapts dynamically to changing real-time conditions.*
Requirement 6: Use of additional storage resources for data replication will be minimized.	✓	Kashya Data Protection Appliance *uses differential approach to minimize the use of additional storage without compromising disaster recovery capabilities.*
Requirement 7: All transfer of data will be secure.	✓	Kashya Data Protection Appliance *protects data from unauthorized access from the time it is sent to the local Kashya appliance to the time it is written to the secondary storage subsystem.*
Requirement 8: Regular testing of DR capabilities at secondary site must be facilitated.	✓	Kashya Data Protection Appliance *enables direct Read/Write access to data at the secondary site, without the need to first make an additional copy. This facilitates regular testing of DR plans at the secondary site, without impacting the replication process.*
Requirement 9: The solution will be universal, scalable and always available.	✓	Kashya Data Protection Appliance *is an end-to-end solution for data replication across heterogeneous server, SAN and storage platforms. The system employs a cluster of multiple active-active independent appliances, resulting in unlimited, non-disruptive scalability and always-on availability*