

Mid-Tier Array Replication — Large Array Features at Mid-Tier Pricing

Analyst: Dianne McAdam

Management Summary

Years ago, there were few options to move data offsite to a safe location. One of these options involved a large truck, lots of tape cartridges and a safe ‘bunker’ or storage facility. After the completion of a backup cycle, the current backup tapes were loaded into a truck to be transported to the bunker. If data needed to be restored, the backup tapes were then loaded into the truck and driven back to the data center. This approach still works today for some applications. However, those applications that must be available all of the time require faster recovery times than trucking allows.

In 1994, EMC delivered the first commercially available, array-based replication product for its high-end storage, called *Symmetrix Remote Data Facility*, or more commonly, *SRDF*. Other vendors of high-end array products, such as HDS, HP, and IBM quickly followed suit. Disk arrays could be located within the same building, campus, or separated by hundreds or thousands of miles. When an application updated the local volume, the remote volume received the update a short time later. If the local disk array fails, then the data can be easily retrieved from the remote controller. Recovery times were dramatically improved by disk replication since bits on a wire can travel at the speed of light, but trucks must obey posted speed limits.

Furthermore, in the 1990s communications lines were expensive. The cost of the high-end systems added to the monthly lease of the communications costs, all of which made these solutions unaffordable to smaller enterprise customers. However, a lot has changed since the mid-1990s. The cost of telecommunications lines has become more affordable and storage vendors have developed replication products for lower cost “mid-tier” disk arrays. These mid-tier replication products may not have all of the bells and whistles of their big brother high-end disk arrays, but they have plenty of features that can support the needs of the smaller enterprises.

The first versions of mid-tier replication products were very limited in capability. Some only supported very limited distances and were only suitable for campus environments. Fortunately, the solutions available today can support many different configurations and distances. If you haven’t taken a look at these products lately, you will be surprised what they can do today. These products deserve a second look for those enterprises that are looking for affordable remote replication solutions. Read on to see what features are available today.

IN THIS ISSUE

➤ Types of Replication.....	2
➤ What Other Features Do You Need?.....	2
➤ Conclusion	4

Types of Replication

There are many different mid-tier array replication products on the market today and no two products are identically featured. They differ in the types of replication that are supported, the number of mirrors that are supported, and the availability of features that can make the life of IT administrators easier. So which features do you need? It depends on your environment.

Short Distance Replication

The first cars built in the United States could be ordered in any color that you wanted, as long as it was black. The same was true of the first mid-tier array-based replication products. You could get them in any option, provided you needed to replicate over short distances.

Synchronous replication ensures that the remote mirror and the local mirror are always exact images of each other. An update to the local mirror is transmitted to the remote mirror and the application is notified that the write operation has completed *after* the remote mirror has been updated. Since the application must wait for the acknowledgement that the update has been completed at the remote site, latency is added to every update - the greater the distance, the greater the latency. Therefore, synchronous replication is best suited for shorter distances, such as replicating within a campus or within several miles. Of course, your mileage will vary based on the traffic on the communications lines and the application's tolerance for latency.

Long Distance Replication

Enterprises that need to replicate across hundreds or thousands of miles had to wait until asynchronous replication became available for mid-tier arrays. Asynchronous replication informs the application that the update is completed *before* the remote mirror is updated. So, there is no latency and replication can span large geographical areas.

Multi-volume Applications

Some applications data can be stored on one volume. Nevertheless, many applications, such as databases, span several different volumes. Multi-volume applications require the *consistency group* feature to ensure that the remote volumes are in a consistent state.

A consistency group is a logical grouping of interrelated volumes. When the link supporting one of the volumes within a consistency group fails, the software suspends replicating the remaining volumes within that group ensuring that all of the related volumes in the remote site remain static from the point of the link failure.

What Other Features Do You Need?

There are other many features available today for mid-tier replication products and vendors continue to enhance their offerings. If you are evaluating array-based replication, the following questions will help to determine which vendors solution will meet your needs.

Synchronous Replication

If you need to replicate within a campus environment, or at distances less than 100 miles, then synchronous replication usually is the best approach.¹

Cost

1. *Is there a separate license charge for synchronous replication? Is it based on the model of the controller? Or is it capacity based? What is the warranty period? What is the cost of maintenance after the warranty period expires?*
 - You may be pleasantly surprised to discover that the cost of many of these replication licenses is rather inexpensive.
2. *Are there any costs associated with assessments? Will the vendor study the current workload to determine the correct configuration? What is the cost of this assessment?*
3. *Are there any costs associated with installation or is this service included in the purchase costs?*

Configuration

4. *Must you replicate between the same models of controllers or can you different models of controllers on either side of the link? Can newer generation of controllers replicate to older generation models? Does the same version of firmware have to run on both controllers?*²
 - If not, it is easier to migrate to newer releases of firmware, since not all arrays have to be upgraded at the same time.
5. *How many arrays can be connected to one array?*

¹ Synchronous replication requires the application to wait until the remote array has confirmed that it has completed writing the new information before proceeding. This takes time and the time is increased by the distance between arrays. It is synchronous because the replication is happening in "real time" at multiple locations.

² There are two issues here. *Can replication be done between dissimilar arrays (possibly from different vendors)? Can different generations of the same array replicate without problem?*

- Some vendors support one array connecting to several different arrays.
6. *How many volumes can be replicated?*
 - This number varies by vendor and by array model.
 7. *All solutions allow a one-to-one mirrored relationship (one volume is mirrored to another volume). Some solutions allow one volume to be mirrored to two different volumes. Do you need to replicate one volume to two different locations?*
 - If so, look for the vendor that supports one-to-two (or more) mirroring.
 8. *What types of links are supported? Most vendors support Fibre Channel and may support Fibre Channel over IP or direct IP connections. How many links can be supported? Are the links bi-directional?*
 - It is always a good practice to configure redundant links.
 9. *What size and number of links are required to support current and future workloads? Most vendors have tools to determine network bandwidth requirements.*
 - With synchronous replication, having insufficient network bandwidth can cause performance problems.
 10. *What are the vendor's distance recommendations?*
 - Some vendors have tested configurations to certain distances using a combination of Fibre Channel and extended-distance communications equipment.
 - Remember that synchronous replication adds latency to applications. Vendors may have specific recommendations for configuring applications, such as *Oracle* or *Microsoft Exchange*.
 11. *How much bandwidth is required?*
 - A workload study should be conducted with your storage and network providers to determine how much bandwidth is required. This study should be performed early to set realistic expectations about the cost of bandwidth and the amount of data that can realistically be mirrored.
 - A general rule of thumb for disk mirroring is that one TB of online transaction data will generate about 1-2 MBs per second of write data. A one TB of batch or sequential workload will generate about 6 -7 MBs per second of write data. This rule of thumb is linear. That is, if you need to mirror 5 TBs of OLTP storage, then you can estimate that the workload will generate about 5 to 10 MBs per second of write data. Of course,

your mileage may vary and an actual workload study *should* be completed.

- Remember to calculate peak workloads. In general, peak workloads are about 2.5 to 3 times higher than average work loads.
- Some inter-SAN connections and channel extenders may support compression, which can reduce the bandwidth required.

Implementation

12. *Must the replication software be installed on each array or is it 'turned on' by a license key? How long will it take to install the software?*
13. *Every implementation requires that the entire volume be replicated to the remote site initially. How long does the vendor predict that this will take? Will the first-time full synchronization affect application performance? Are tools available to monitor the status of the synchronization? Can the synchronization priority be modified? If so, how granular is the priority scheme? Can priorities be modified dynamically to slow down the replication process during periods of heavy application processing?*

Operations

14. *After a link failure has been resolved, do you have to copy the entire contents of the local volume to the remote volume to get the two volumes 'in sync'? Or does the software keep track of the changes that have occurred during the failure and only copies the delta changes?*
15. *Can the role of the local and remote mirror be reversed; that is, can you reverse the mirroring operation from the local array to the remote array?*
 - This feature is important when restoring the contents of the local volume after a failure of the local array. This feature is sometimes called "role reversal" or "personality swap". After the failure has been rectified, must you restore the entire contents of the local volume or can the software send over only the delta changes that have occurred since the failure?
16. *Can you access the remote mirror?*
 - You may want to access the remote mirror in read only mode to perform other tasks, such as running backups or loading data warehouses.
17. *If you cannot access the remote mirror directly, can you make a copy of it to perform these additional tasks?*
 - Some vendors allow administrators to create snapshots of remote mirrors.

18. *Is the snapshot software integrated with the remote mirroring software? If not, can you easily create scripts to create snapshots?*
19. *If you need to create snapshots, is there an additional licensing fee required?*
20. *Are consistency groups supported for multi-volume applications? If so, how many consistency groups can be created on one array? How many volumes can exist within a consistency group?*

Management

21. *What management software is available to manage the replication? What is the cost of this software? Can the software manage both the local and remote mirrors from one console? Can the software manage multiple local and remote arrays? How easy is it to set up the initial mirrored pairs? Is training required to run the management software or is it easy to use?*

Asynchronous Replication

If you need to replicate to other site that is hundreds or thousands of miles away, then asynchronous replication is your only choice. There are additional questions that need to be answered for asynchronous replication.

Cost

22. *Is there a separate license charge for asynchronous replication? Is it based on the model of the array? Or is it capacity based? What is the warranty period?*
 - Some vendors include synchronous and asynchronous replication in one package, while others price it separately.

Configuration

23. *What devices have been certified by the vendor to work with their firmware to extend communications lines?*
24. *How are updates sent over the link? Some vendors have implemented asynchronous replication to send updates over the link as they occur; others send over updates periodically in batches. If the vendor sends batches, do you have any controls to determine how often the batches are replicated?*
25. *How much bandwidth is required? It is critical to have sufficient bandwidth to support synchronous replication activities since the application must wait until the update has been applied to the remote disk system.*
 - Having sufficient bandwidth is also important for asynchronous replication operations. Asynchronous replication informs the application that the I/O is completed before the remote mirror is updated. However, if there

is not enough bandwidth available, then the number of updates that need to be sent to the remote mirror will ‘stack up’. This can result in the remote mirror being seconds or minutes ‘behind’ the local mirror. Recovery Point Objective (RPO) defines the point in time in which the applications can be restored if a failure occurs. With synchronous replication, the RPO is zero – that is, the local and remote mirror are identical and there is no difference in time between the two volumes. If sufficient bandwidth is available for asynchronous replication products that support continuous updating, then the RPO approaches zero.

26. *Is write ordering maintained?*


- Some provide write consistency; that is, they sequence the updates to ensure that they are received at the remote array in the same order that they were received locally. If write ordering is not maintained, then there is a possibility that updates will be applied to remote volumes in the wrong order, which can result in inconsistent volumes after a failure. *If the vendor does not support write ordering, do they have other techniques to ensure the remote volumes are consistent after a failure? If not, what procedures do you have to set up to establish consistent points periodically?*

Operations

27. *Are consistency groups supported? If so, how many and how many volumes can exist within a group?*
28. *Can you switch between synchronous and asynchronous replication?*
 - Synchronous replication can affect application performance. Switching to asynchronous mode during heavy bandwidth utilization can improve application performance.

Conclusion

Array-based replication allows enterprises to mirror data to remote locations without using server cycles. For many years, these products were only available on high-end arrays. The initial replication products for mid-tier arrays were limited in distance and in features. But that has changed. Today’s mid-tier array replication products have many of the features formerly only supported on more costly arrays at mid-tier pricing. That makes these products an attractive option, worthy of a fresh look.



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

Dianne McAdam is Director of Enterprise Information Assurance for the Clipper Group. She brings over three decades of experience as a data center director, educator, technical programmer, systems engineer, and manager for industry-leading vendors. Dianne has held the position of senior analyst at Data Mobility Group and at Illuminata. Before that, she was a technical presentation specialist at EMC's Executive Briefing Center. At Hitachi Data Systems, she served as performance and capacity planning systems engineer and as a systems engineering manager. She also worked at StorageTek as a virtual tape and disk specialist; at Sun Microsystems, as an enterprise storage specialist; and at several large corporations as technical services directors. Dianne earned a Bachelor's and Master's degree in mathematics from Hofstra University in New York.

- ***Reach Dianne McAdam via e-mail at dianne.mcadam@clipper.com or at 781-235-0085 Ext. 212. (Please dial "212" 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, Clipper Notes, 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.