IBM eserver

Geographic Mirroring in the SAP System Environment on eServer iSeries

Dennis Bourget
SAP on iSeries Development Support
December 2004

Table of Contents

Introduction

This paper discusses the various influences geographic mirroring may have on the SAP system environment running on the IBM eServer iSeries. Topics discussed include network influences, the general behavior of geographic mirroring, CPU/memory considerations, switchover times, and the general performance of an SAP R/3 Enterprise system when running in a geographically mirrored environment.

What Is Geographic Mirroring?

Geographic mirroring is a subfunction of cross-site mirroring (XSM) that generates a mirror image of an independent ASP (iASP) on a system that is (optionally) geographically distant from the originating site for availability or protection purposes.

An iASP contains objects, the directories and libraries that contain the objects, and other object attributes, such as object authorizations, in a clustering environment.

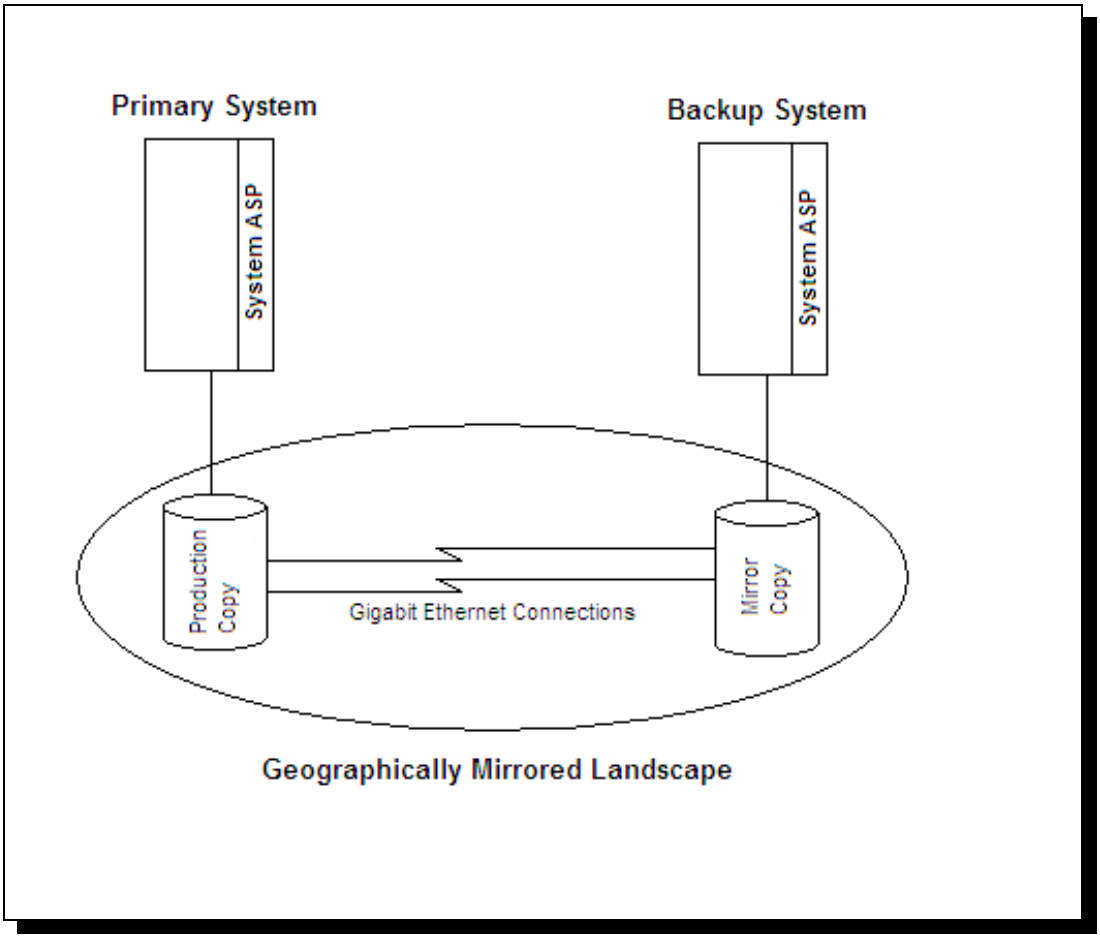

How Geographic Mirroring Works

Geographic mirroring executes in the storage management tasks in the SLIC layer of the iSeries. The mirroring is carried out by sending data segments from the primary system containing the original server data, via a network connection, to a backup system that owns the mirror copy. The backup system is then responsible for writing this data segment to the mirror copy. These disk segments are independent of the type of data on the independent ASP. Hence, it is not the number of objects on the independent ASP that influence the synchronization performance of the geographically mirrored iASPs, but rather the total amount of data stored on the iASP.

Synchronous vs. Asynchronous Performance Mode of Geographically Mirrored Independent ASPs

When configuring geographic mirroring, the user has the option of performing synchronous updates, or asynchronous updates to the mirror copy. Both options ensure that the mirror copy is always usable (eligible to become a production copy).

Synchronous updates force the task performing the write operation on the primary system to wait for the information being written to disk to be sent to the backup system and written to the mirror copy. Once it is written on both the production copy on the primary system and the mirror copy on the backup system, the task on the primary system is allowed to continue.

Asynchronous updates work in the same manner as the synchronous updates, but rather than waiting for the write on the backup system to complete, the task performing the write on the primary system only waits for the information to be sent to the backup system. Once the information is in the backup system's memory, the task performing the write on the primary system is allowed to continue. This eliminates the need for the primary system to wait for the data to be written to the mirror copy.

Synchronization

After the clustered environment with geographic mirroring is configured, the production copy and the mirror copy must perform an initial synchronization.

Re-synchronizaion is required when the production copy on the primary system and its mirror copy on the backup system become out of sync. This can be caused by various situations, including hardware failures, network failures, or scheduled system maintenance. When this occurs, re-synchronization must perform the same steps as an initial synchronization, and copy all the data from the production copy to the mirror copy.

During both the initial synchronization and any re-synchronization, the iASP is still available to the production system, but the mirror copy on the backup system is unusable (not eligible to become the production copy) until the synchronization has completed.

The amount of time needed for synchronization can be influenced by the amount of activity on the iASP being synchronized. The lesser the activity on the independent ASP, the sooner the synchronization will complete. If there are any write requests to the portion of the iASP that has already been synchronized, these write requests will have priority over the un-synchronized portion of the iASP, and will be sent to the mirror copy first. This helps to ensure the previously synchronized portion of the iASP is kept current.

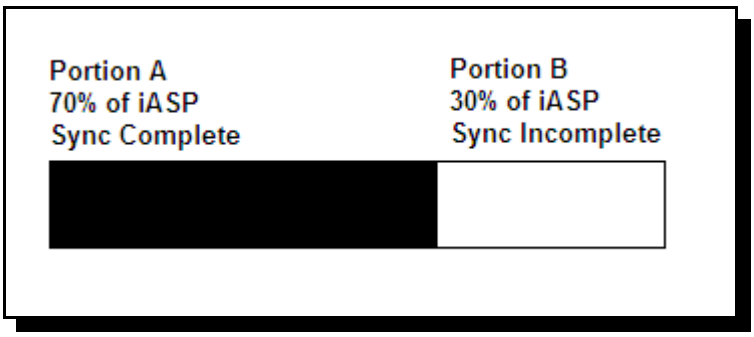

As the synchronized portion of the independent ASP (Portion A) grows, so does the likelihood of a write being performed on this portion of the iASP. Each of these write requests will cause the remaining synchronization (Portion B) to be delayed. As the number of writes on the system grows, so will the amount of time required for the synchronization to complete.

The effect of the re-synchronization on a production system is affected by the “Resume Priority” property of geographic mirroring. Three options are available: Low; Medium; and High. As the priority increases, so does the amount of CPU used by the production system when performing the synchronization/re-synchronization. While more CPU may be used, the higher the priority, the sooner the re-synchronization will complete.

Switchover Scenarios and Behavior

A switchover occurs when the mirror copy assumes the role of the production copy. The original production copy then assumes the role of the mirror copy, or is unavailable due to various reasons, including system failure, network failure, or scheduled system maintenance. In the case of another node at the production site that is configured as the next backup, then the switchover moves the production copy to that node, and the mirror copy remains the mirror copy.

There are two scenarios that can cause the production copy to change to the next backup node in the cluster, or, dependent on how the cluster is configured, to have the mirror copy assume the role of the production copy. The first scenario is a switchover. A switchover is initiated by the user using the CHGCRGPRI command, or via the iSeries Navigator.

The second scenario is a failover. A failover is automatic when the Cluster Resource Group (CRG) is active.

The iASPs may be online or off-line when the switchover, or failover, occurs. A cluster node in "partition" status does not cause failover. The user must explicitly change the node status from "partition" to "failed" to cause failover.

The time required for the switchover consists of the time required to perform the following steps: (**Note**: this is a description of a controlled switchover, when both systems are available)

1. Verification that the cluster group is able to be switched.
2. The production iASP is moved into an unavailable state. When this occurs, any jobs associated with it are also ended.
3. The backup iASP is brought into an available state. At this point the iASPs have switched roles, and the former mirror copy assumes the role of the production copy.

The amount of time required for a switchover is dependent upon the state of the independent ASP at the time the switchover occurs.

Performance studies on an independent iASP containing one Terabyte of data have shown that both planned and unplanned switchover times range from 6 to 10 minutes, with the planned switchovers finishing sooner than unplanned switchovers.

General Performance Considerations

There are several factors that need to be considered when configuring your system for Geographic Mirroring.

The first consideration is the network configuration. The faster the network connection, and the more bandwidth available, the faster the synchronization updates to the mirror copy are completed. While only one network connection is required for geographic mirroring, the current recommendation is to have at least two to four gigabit ethernet lines dedicated to geographic mirroring. Synchronization and update performance does not increase linearly as more lines are added. In most cases, two dedicated gigabit ethernet lines provide the most benefit to performance. Having a third or fourth line provides some performance benefit. By dedicating these addition ethernet lines to geographic mirroring, we help ensure that the geographic mirroring is less likely to be interrupted if one or more network connections are lost. The remaining connections will continue the geographic mirroring between the production copy and the mirror copy.

Another important consideration of Geographical Mirroring involves the User IDs and Group IDs of users owning objects on the iASP. It is important to ensure that the owners of the objects on the backup system have the same ID values as on the production system. Leaving these values different can have a dramatic effect on switchover times. The normal switchover time for a iASP with one Terabyte of data is approximately 6 to 10 minutes on an inactive system. In our test environment, with only a 90 Gigabyte library, it took in excess of 90 minutes for the switchover to complete, when the User IDs and Group IDs were different.

CPU and Memory Considerations

The impact of synchronization on CPU usage is dependent upon how geographic mirroring is configured, as well as the speed of the network being used for synchronization. In a controlled test environment, using a geographically mirrored iASP with one Terabyte of data, sychronizations using four 1-gigabit ethernet connections would finish in approximately 63 to 96 minutes. The synchronization used an additional 9% to 16% of the CPU on the production system, and an additional 12% to 18% of the CPU on the backup system. The level at which the "Resume Priority" is set can have an influence on the amount of CPU utilization, and total time required for synchronization. When the test environments had a "Resume Priority" setting of "High", it used 16% of the CPU, but the synchronization completed in 63 minutes. When the "Resume Priority was set to "Low", the synchronization used only 9% of the CPU, but the synchronization took 96 minutes to complete. Hence, a higher "Resume Priority" will use more CPU, but synchronization will complete sooner, and a lower "Resume Priority" will use less CPU, but synchronization will take longer to complete.

Currently, there is no significant impact on the amount of memory used on the system, and no additional memory is expected to be needed on systems using geographic mirroring. There will be a small overhead due to more storage management tasks running in the background, but the footprint required by these is minimal.

While doing normal updates, an additional 15% of CPU usage is expected on a system containing a geographically mirrored iASP, as compared to a system containing an iASP that is not using geographic mirroring.

Performance Observations of an SAP R/3 Enterprise Environment

In our test environment, we installed an SAP R/3 Enterprise system into a clustered environment, using a geographically mirrored iASP to contain the SAP database library and its journal.

Our production system was setup as a second partition on a 890 iSeries Model, using 4 CPU, and 8 Gigabyte of memory.

The backup system was a 520 iSeries, using 1 CPU, and 8 Gigabyte of memory.

Both systems had a system ASP with one Terabyte of DASD, and a geographically mirrored iASP with one Terabyte of configured DASD. Of the one Terabyte in the iASP, only 20% of the iASP was used. Both systems had i5/OS version V5R3 installed.

In order to observe how an SAP workload may be affected, we performed a client copy (transaction SCCL) in a geographically mirrored environment with three different network configurations. The first network configuration consisted of the recommended two gigabit lines dedicated to geographic mirroring, the second consisted of single gigabit line, the third configuration used a 100 megabit line. We then compared the performance of the client copy transaction in these three environments to an environment consisting of the production system with iASP configured, but no geographic mirroring.

When comparing the client copies of both a single gigabit network connection, and a two gigabit line to the client copies run on an iASP with no geographic mirroring, we found no degradation in total run time. This indicates that both gigabit line configurations provided a sufficient amount of bandwidth to accommodate the updates performed by a client copy.

The third environment, with only one 100 megabit connection, is substantially less bandwidth than is recommended, and should not be used if possible. In a normal environment, the geographic mirroring should be spread amongst several network connections. This will increase bandwidth, allowing updates and synchronization to complete sooner, as well as providing a more reliable geographically mirrored environment. When comparing a client copy in this environment to the client copies run on an iASP with no geographic mirroring, the client copy took 20% longer to complete in the 100mb geographically mirrored environment.

Conclusion

When installing and running a SAP R/3 Enterprise system in an iSeries environment with geographically mirrored iASPs, the following must be considered when configuring the environment.

- The speed and amount of bandwidth available to Geographic Mirroring can have a substantial impact on the performance of synchronization and updates. If performance begins to degrade in a Geographically Mirrored environment, causes may be due to a slowdown in the network used, or an increase in the amount of data being sent.

- Having multiple network connections and paths between the production iASP and the Geographically Mirrored iASP decreases the possibility of Geographic Mirroring being interrupted if a single connection fails.

- Ensure that the User IDs and Group IDs are identical for all users storing data on the geographically mirrored iASP. Unmatched IDs can result in a 15 fold increase in the amount of time required to perform a switchover.

- Geographic mirroring can add an additional 15% of CPU usage onto a system.

- Geographic mirroring has no effect on read operations. Since all data being read is on the production iASP, there is no need to access the network connection to the mirror copy. Due to this, a workload that is more read intensive will see less impact from geographic mirroring, when compared to an environment that is more write intensive that must send updates across the network.

Related Documents/Publications

- *Clustering and IASPs for Higher Availability on the IBM eServer iSeries Server*, SG24-5194
- *Independent ASP Performance Study on the IBM eServer iSeries Server*, REDP-3771
- *iSeries Performance Capabilities Reference i5/OS Version 5, Release 3*, SC41-0607

For more information on how to set up the SAP system environment on iASPs, see SAP Note 568820 in the Online Service System.

Trademarks

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.
The following terms are trademarks or registered trademarks of the IBM Corporation in the United States or other countries or both:

e(logo)server	IBM	iSeries
eServer	IBM (logo)	i5/OS

SAP, SAP R/3, and SAP R/3 Enterprise are the trademarks or registered trademarks of SAP AG in Germany and in several other countries.

Disclaimer - System Performance

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Information is provided "AS IS" without warranty of any kind. Mention or reference to non-IBM products is for informational purposes only and does not constitute an endorsement of such products by IBM.