

Auto-Replication: Automated replication for Failover Clusters

Swapnil Mirjolkar^[1], Ankur Mourya^[2], Vinay Karajagi^[3] Prof. Dashrath Mane^[4]

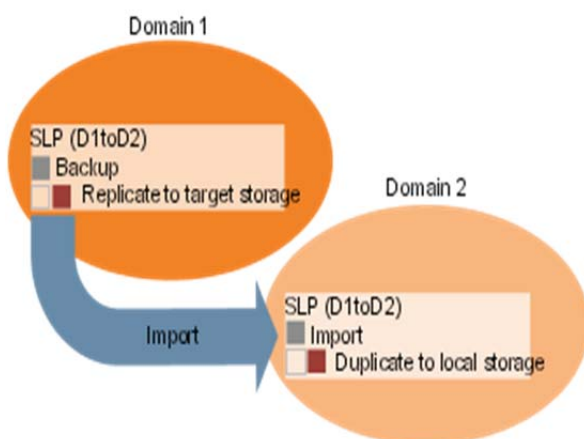
VES Institute Of Technology (VESIT),
Chembur Mumbai:-74, Maharashtra, India

Abstract: Auto Replication is widely implied concept mainly in network field pertaining to domain controllers, database servers, high-availability application servers and so on. It focuses on controlled redundancy of sensitive and specific data over a distributed network. Auto Replication implemented on Dell EqualLogic PS Series network switches has been selected as a muse to better understand this concept.

1. INTRODUCTION

A replica represents the contents of a volume at the time the replica was created. Each replicated volume has a replica set, which is the set of replicas created over time. A volume and its replicas are always stored on different groups connected by a robust network link. Separating the groups geographically protects volume data in the event of a complete site disaster. To replicate volume data, you configure the group that contains the volume and the group that will store the replicas as replication partners. A group can have multiple partners, but a volume can replicate to only one partner. Each partner plays a role in the replication of a volume, and you can monitor replication activity from either partner:

- **Primary group.** This is the group storing the volume. The primary group administrator configures the secondary group as a replication partner and initiates the replication operation. Replication from the primary group is considered outbound.
- **Secondary group.** This is the group storing the replicas. The secondary group administrator configures the primary group as a replication partner and provides space for the replicas. Replication to the secondary group is considered inbound. Mutual authentication using passwords provides security between partners.



2. OVERVIEW STEPS

1. Plan the replication configuration for your groups and volumes. For each volume, determine the replication space requirements for the primary and secondary groups. This will help you determine whether you have the local space required for replication and also choose the best replication partner for each volume.
2. Configure one or more replication partners. For each partner that you configure, you must log into the partner and configure the primary group as a replication partner, making sure you delegate sufficient space to the primary group.
3. Configure each volume for replication. Specify the replication partner, local replication reserve, replica reserve on the partner, and whether to maintain the failback snapshot.
4. Create a replica on demand, or set up a schedule to create replicas automatically at a specified time and frequency. Using a schedule, you can specify how many replicas to keep.
5. Periodically check to make sure that the replicas have completed and that the correct number of replicas is stored on the partner.
6. Adjust values over time as needed. For example:
 - Adjust the schedule to create or keep more or fewer replicas, or to create them at a different interval or time.
 - Increase or decrease the amount of delegated space on the partner, depending on your actual space usage and needs.
 - Increase or decrease any volume's replica reserve.

3. PLANNING CONSIDERATIONS

Answers to the following questions will help you determine which groups to configure as replication partners, how often to replicate the volumes, how many replicas to keep, how to coordinate your replication plan, and whether or not to enable failback on the volumes:

- How large are the volumes?
- How much and how often does volume data change?
- How many copies of the volume (replicas) do you want to keep?
- How much space is available on each group to store a partner's replicas?
- What are your business recovery needs for the applications using the volumes?

o **Recovery Time Objective (RTO)** —The amount of time within which a business process must be restored after a disaster or disruption. Your RTO determines whether or not to keep a failback snapshot.

o **Recovery Point Objective (RPO)** —The acceptable amount of data loss for a specific amount of time; in other words, the currency or recency of data in the replicas. Your RPO determines the necessary frequency of replication.

4. SPACE CONSIDERATION

Partners: A PS Series group can have up to 16 replication partners¹. This gives you significant flexibility in choosing a replication configuration to meet your needs, for example:

- One group replicates to another group that does not replicate its own volumes. The limits mentioned in this Technical Report apply to PS Series groups that contain at least one member that is not a PS4000. In a group of only PS4000 arrays, the limits are two members per group, two replication partners, and 32 volumes configured for replication.
- Two groups replicate to each other.
- One group replicates different volumes to different groups.
- Several groups replicate to the same secondary group.

Local Space: To replicate, volumes need space on their own group. The local space, called **local replication reserve**, keeps track of changes that occur to the volume while it is being replicated. It also stores a failback snapshot, if you choose to keep one. Because replicas and failback snapshots are a representation of the volume data at a precise point in time, the group must maintain that data point until the replication completes. To allow the volume to remain online and in use during this time, the group uses the local replication reserve to temporarily track changes to the volume while the operation is underway. When the replication is complete, the local replication reserve space is freed, unless you configured the volume to keep a failback snapshot.

Space on a Partner: Volumes need space on the partner group to store their replicas. The amount of space set aside for all the replicas from a partner is called the **delegated space**. Groups can have different amounts of delegated space for each of their partners. The portion of delegated space that is set aside for each volume is called the **replica reserve**. Administrators can increase a partner's delegated space as needed, as long as additional free space exists, or decrease it, as long as no replicas will be deleted. A volume can have only one replication partner. However, different volumes on the same series group can replicate to different partners. For example: on Group One, the replication partners are Group Two and Group Three. On Group One, the volumes and their replication partners can be set up as follows: Vol1 GroupTwo, Vol2 and Vol3 - GroupThree

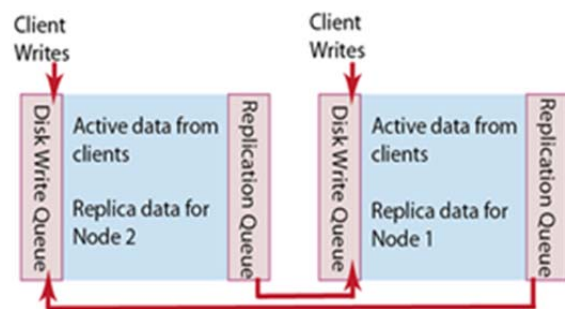
Volume Collections: A volume collection is a group of up to eight volumes that you want to manage together, based on your business needs. For example, a volume collection can consist of a volume for a database and the volume for its log file; volumes that store electronic mail; volumes for a particular group of users or a department; or any other volumes that you want to replicate at the same time. All the

volumes in a collection must be configured to replicate to the same partner.

Replicas: Each volume can have up to 512 replicas stored on its partner. Based on your RTO and RPO, you can specify how often to replicate a volume and how many replicas you want to keep on the partner group.

For example, if your volume data does not often change, you might decide to replicate the volume once per week. Depending on how much data is replicated each time and the amount of replica reserve, you could keep many weeks' or months' worth of replicas on the partner. On the other hand, if a volume's data changes frequently and you want to minimize the amount of potential transaction loss in the event of a disaster, you might decide to replicate the volume daily or even hourly. You could keep a full day's worth of replicas on the partner (24 replicas, if you replicate the volume hourly). The replication frequency and the number of replicas to keep depend on your business needs, and are limited only by the configuration limits and the amount of space available on the volume's replication partner. Use the following formula to determine the number of replicas to keep: Frequency of replications x time period to retain = Number of replicas to keep

For example: 2 replications per day x 14 days = 28 replicas



Schedules: Using schedules can help ensure your volumes and volume collections are replicated at the best intervals and times to meet your business needs.

For example, consider a volume that rarely changes. You can configure a schedule that replicates its data once a week on Sunday mornings at 2:00 AM. If you decide to increase or decrease the replication frequency, you can modify the schedule at any time. To save time, you can create one schedule and then re-use it for many volumes or collections.

5. SECONDARY GROUP CONSIDERATION

Because replication involves two groups, the people who administer the groups must coordinate their efforts. The same person can log into both groups and configure replication and manage all the operations alone, but more often, it will be a joint effort between two or more locations and administrators. To configure groups as replication partners, each administrator needs the following information:

- Group name and IP address of the partner
- Amount of space each group will delegate to the other, if any
- Mutual authentication passwords that enable the two groups to communicate with each other

6. FAILOVER AND FAILBACK CONSIDERATION

In the context of EqualLogic PS Series storage, failover is the mechanism by which you can temporarily host a volume on the secondary group and move its users from the primary group to the secondary group in the event of a disaster on the primary group. At some point, you can also fail back--return the volume and its users to the original primary group. Configuring the volume to keep a failback snapshot can help minimize the amount of data that must be replicated back to the primary group when you are ready to restore the original configuration.

Without a failback snapshot, you can still return to hosting the volume on its original group, but you must replicate the entire volume's contents, which could take hours or even days depending on the size of the communications link and the amount of data to transmit. If you decide to enable failback support, you can specify it at the time you initially configure replication on the volume, or at any time thereafter.

Note: You must create a replica to establish the failback snapshot before you attempt to fail back. Because the failback snapshot uses space from the volume's local replication reserve, you may want to increase that value to make sure replication continues to proceed successfully.

7. CASE STUDY WITH EQUAL LOGIC

The EqualLogic PS Series builds on a unique peer-storage architecture that is designed to provide the ability to spread the load across multiple array members to provide a SAN solution that scales with the customer's needs. This "pay as you grow" model, allows customers to add arrays as their business demands increase the need for more storage capacity or more I/O capacity. Every EqualLogic array includes additional enterprise class features such as snapshots, clones, replication and all-inclusive software: Group Manager, SAN Headquarters (SANHQ) and Host Integration Toolkits. The built-in snapshot feature enables quick recovery of files, clones for recovery of files or volumes, and the replication feature allows the implementation of disaster recovery initiatives. The EqualLogic software includes storage system-based replication. This feature (called Auto-Replication) provides the ability to replicate data volumes to peer EqualLogic storage arrays situated in remote locations without setting the volumes offline. Auto-Replication provides a disaster-recovery option in case the original volume (or the entire EqualLogic group) is destroyed or otherwise becomes unavailable. Auto-Replication is a point-in-time replication solution that offers extended-distance replication. EqualLogic Auto-Replication provides asynchronous, incremental data synchronization between primary and secondary replicas. Scheduled replication events update the remote copy of the data with all the changes that occurred on the primary copy since the last replication event occurred.

Key benefits of using EqualLogic Auto-Replication:

Using Auto-Replication with an EqualLogic SAN as part of your storage infrastructure can provide multiple benefits:

SAN-based replication is included at no additional cost: EqualLogic Auto-Replication replicates volumes to remote

sites over any distance by leveraging existing IP network infrastructure.

Ease of use: EqualLogic Group Manager and SAN Headquarters (SANHQ) are included at no additional cost. These tools provide easy-to-use GUI based tools for managing and monitoring EqualLogic arrays, including all replication oriented tasks.

Manual Transfer Utility: The Manual Transfer Utility is also included at no additional cost. This host based tool integrates with the native replication function of the EqualLogic Firmware to provide secure transfer of large amounts of data between EqualLogic groups using removable media. The Manual Transfer Utility is beneficial in environments where data protection is critical but bandwidth is limited.

Auto-Replication addresses varying needs using powerful features and configuration flexibility:

- Multiple recovery points are efficiently stored.
- Per-volume replication schedules permit varying service levels.
- You can fast failback to the primary site by synchronizing only the data that has changed while the secondary site was in use.
- One-way, reciprocal, one-to-many or many-to-one replication paths are possible.
- Thin replicas provide space efficiency.

EqualLogic Auto-Replication is used to replicate volumes between different groups as way to protect against data loss. The two groups must be connected through a TCP/IP based network. This means that the physical distance between the groups is not limited. The replication partner group can be located in the same datacenter, or it can be in a remote location. In practice, the actual bandwidth and latency characteristics of the network connection between replication groups must be able to support the amount of data that needs to be replicated and the time window in which replication needs to occur. IT administrators can enable replication through the Group Manager GUI or the Group Manager CLI. Once enabled, volume replication functions can be managed using the GUI, the CLI, or Auto-Snapshot Manager (ASM) tools. ASM tools are included as part of the Host Integration Toolkits for Microsoft and VMware. A replica can be created from a single volume or a volume collection. A volume collection allows up to eight volumes to be grouped together so that the replication process for that collection can be managed as if it was a single volume. This is useful when replicating multiple volumes that belong to the same application.

Auto-Replication initially creates a copy on a secondary storage system, and then synchronizes the changed data to the replica copy. A replica represents the contents of a volume at a point in time at which the replica was created. This type of replication is often referred to as "point-in-time replication" because the replica copies the state of the volume at time the replication is initiated. The frequency in which replication occurs determines how old the replica becomes relative to the current state of the source volume.

8. CONCLUSION

This can be described as only a brief description of what exactly is the working style for Auto-Replication module. With each day new systems and working algorithms are being innovated and implemented for increased efficiency in Distributed systems and management.

REFERENCES

1. IEEE Transactions on Reliability, VOL 47, NO 1, 1998 March 'Module Replication for Fault-Tolerant Real-Time Distributed Systems' by Theodora A. Varvarigou and John Trotter
2. Issue featured in Summer USENIX Conference, Anaheim. CA, June 1990 'Implementation of the Ficus Replicated File System' by Richard G. Guy and John S. Heidemann
3. www.equallogic.com/resourcecenter/assetview.aspx?id=5337
4. www.equallogic.com/resourcecenter/assetview.aspx?id=9723
5. http://www.cisco.com/c/en/us/td/docs/security/asa/asa91/configuration/general/asa_91_general_config/ha_failover.html#pgfId-1160857
6. <http://en.wikipedia.org/wiki/Failover>