

Data Replication Cookbook

A Hypothetical Look at Replication Planning

TABLE OF CONTENTS

Replication Planning	3
The Environment	3
A Replication Cookbook	4
Running vxstat	4
Sizing the SRL	5
Sizing Bandwidth	5
Using VERITAS VRadvisor	5
Analyzing Replicated Data	5
Additional Files to Replicate	6
Summary	6
Appendix	7
VERITA Volume Replicator Advisor (VRadvisor)	7

Replication Planning

It is easy to forget when planning for Disaster Recovery that replicated data is only one piece of the solution. The purpose of replication is to provide business continuity in the event of an outage that disrupts access to the organization's main computing infrastructure. This may sound like common sense but it is all too easy to get bogged down in the minutiae of data replication technology and lose sight of the main goal. For example, just because a one TB storage array is dedicated to production systems does not mean that one TB of data must be replicated on a regular basis. Focusing on the one TB of data and the enormous quantities of bandwidth and storage needed to accommodate the replication task distracts from the real reason for replication: to return business systems to normal functioning as quickly as possible. With this in mind it is worth reviewing what data is needed for business systems to operate. Perhaps all one TB of data does not need to be replicated, and, more importantly, small amounts of additional critical data may exist in locations other than the one TB array.

Replication planning is all about identifying exactly what data needs to be replicated, and then ensuring that it is copied to a remote location as quickly and efficiently as possible. This paper follows the analysis of a storage administrator planning to replicate data to a remote data center with VERITAS Volume Manager™ (VxVM). Because the administrator knows that she will have to perform replication planning again for other business applications, she has prudently documented her analysis in a cookbook so that the same steps can be followed next time.

In this scenario, the administrator is faced with a fairly common environment containing Web, application, and database layers. The database in this example is Oracle but could just as easily be IBM DB2 UDB, Microsoft SQL Server, Sybase, or MySQL. The database, however, is just one component in the overall DR solution. This demonstration is to highlight the importance of only replicating what is needed. It also provides some details on common data items that should not be replicated.

The Environment

In this hypothetical configuration the primary data center site is in Boston, with a secondary site in Chicago. The organization has a recovery point objective (RPO) of one hour and a recovery time objective (RTO) of four hours. The RPO and RTO are outputs from a business impact analysis (BIA) process. The analysis determines how much data an application can risk losing before business begins to suffer losses - the RPO - and how quickly an application must be resurrected after a failure in order to avoid business losses - the RTO.

Both data centers employ Sun servers, running Solaris 9. The database system is to be replicated to the secondary data center, where it will also be used to perform off-host backup.

The database system is configured with one TB of storage shared between eight volumes. Figure 1 shows the storage mount points on the Sun Server. This storage is used solely for database data. Operating system and administrative files are stored elsewhere.

```

/apps/Oracle/u001 -- 100GB - Oracle Binaries and dbf files
/apps/Oracle/u002 -- 100GB - dbf files
/apps/Oracle/u003 -- 100GB - dbf files
/apps/Oracle/u004 -- 100GB - dbf files
/apps/Oracle/u005 -- 100GB - dbf files
/apps/Oracle/Backup -- 250GB - dbf files
/apps/Oracle/TempFiles -- 100GB
/apps/Oracle/Archive -- 150GB

```

Figure 1: Storage mount points on the Sun Server

The hardware configuration of each server is not significant for replication considerations, other than to note that each must have at least two host bus adapters (HBA) to support dynamic multipathing (DMP), avoiding a single point of failure, and each should provide sufficient CPU power to run the database.

Performing data replication does not add a noticeable amount of overhead on the database system. Write activity will be doubled, because the system has to write to the data volumes and the VERITAS Volume Replicator™ (VxVR) storage replicator log (SRL).

The SRL receives copies of all replicated I/Os to guarantee consistency of the replicated data. If a network outage occurs between the source and target of the replication, the SRL continues to be written locally and then populates the remote volumes with changed data once connectivity has been restored. With two 100 MBps HBAs the Sun servers are capable of accommodating 17,280 GB of throughput per day – 200 MBps times 86,400 seconds. With this amount of capacity, doubling the number of I/O writes will have no noticeable impact on the system.

A Replication Cookbook

The first step in the process of planning for replication is to determine what is happening in the environment. This means gathering data about the storage system, and how it is being used, on a regular basis.

Running vxstat

Each of the volumes listed in Figure 1 is needed to run the database. To determine exactly how the database volumes are being used, the administrator runs a vxstat.

Running vxstat produces a large amount of data, but will provide essential input to the replication planning process. Vxstat will show the activity on the storage volumes and how much data changes. Figure 2 shows the format of the vxstat command.

```
# vxstat -g oradg -i 120 >> /tmp/vxstat.oradg
```

Figure 2: The vxstat command format.

To obtain a thorough view of the system, vxstat is run for four weeks, and preferably left to run through an end of the quarter period. The amount of time and the period in which the statistics are collected is very important. Systems have different loads at different times, and the objective of this analysis is to design a replication solution that will address the system's load in every period, especially the busy ones.

Sizing the SRL

After examining the results of the vxstat command, the storage administrator determines that there are about 250 GB of changes to data on the database volumes per day. To accommodate a network outage of two days, the VxVR SRL must be 500 GB.

Sizing Bandwidth

Replication requires substantial network bandwidth. Before assessing the amount of bandwidth needed, storage rates, measured in MBps, must be converted to network transfer rates, measured in megabits per second (Mbps). To convert MBps into Mbps, divide MBps by eight ($\text{Mbps} = \text{MBps} / 8$). A network connection supporting a T3 transmission rate is capable of transferring data at 45 Mbps. This is equivalent to 5.625 MBps ($45 \text{ Mbps} / 8$). Therefore, a T3 is capable of supporting 486 GB of data replication per day ($5.625 \text{ MB/s} * 86,400 \text{ seconds}$).

A single T3 connection between the local and remote data centers will be sufficient to accommodate the estimated 250 GB of data changes produced daily by the database system.

Using VERITAS VRadvisor

VERITAS VRadvisor can help to analyze the copious amounts of data gathered during a replication planning exercise. VRadvisor is included in the VERITAS File System™ (VxFS) distribution and can be located in the directory indicated in Figure 3.

```
sf.4.0.sol/windows/VRTSvradv.msi
```

Figure 3: Location of VRadvisor on Windows systems

A data file generated on a UNIX system can be ftp'd or copied to the Windows environment and used as the input to VRadvisor. See the appendix for more information.

Analyzing Replicated Data

When reviewing mounted storage volumes the administrator notices that one volume is named backup. After talking with the database administrator (DBA), it is clear that this volume is used to make a periodic backup of the database to disk. The backup data does not need to be replicated, but the volume must exist at the secondary data center in order for the database to operate successfully. The backup volume can be excluded from the VxVR replicator volume group (RVG). The backup volume handles 50 GB of data changes per day. Eliminating this volume from the RVG reduces the total amount of data being replicated to 200 GB per day.

Further analysis of the mounted storage shows a volume named Archive. After talking to the DBA once more it is clear that this volume is used to store archive logs. Although archive log data is essential to the database system it is not needed to start the database at the secondary site, or to run business applications immediately after a disaster event.

The Archive volumes are responsible for 20 GB of data changes per day – 10 GB for storing the database archive logs and another 10 GB of activity when the logs are compressed. Although the archive logs are important when recovering the database from tape at the disaster recovery site — the logs allow changes, made after the backup tapes were created, to be reapplied — if bandwidth is an issue the log volume can be removed from the RVG, bringing the total amount of data to be replicated down to 180 GB per day

A final review of the database storage identifies another potential candidate for removal from the RVG. The TempFiles volume is used when the database performs large SQL queries. Temporary files are generated by the database engine to store intermediate query results and to maintain the state of the query. These files generate significant amounts of I/O activity and DBAs frequently move temp files to a single storage location to avoid I/O contention with data and index files. Analysis of the vxstat results shows that the temp files for the database are responsible for 120 GB of write I/O activity per day. By ensuring all temp files are stored on a single volume, and then removing this from the RVG, the total amount of data changes replicated per day is reduced to 60 GB.

The review of replicated volumes not only reduces the amount of data being moved across the network, giving an average data rate of 5.5 MBps, but also allows a resizing of the VxVM SRL. A 120 GB SRL will now accommodate sufficient data changes to encompass a two-day network outage.

Additional Files to Replicate

Although the majority of data necessary for successful daily operation of the database is stored on the mounted volumes identified in Figure 1, it is essential to identify and review all sources of data used by the database. Control and configuration files, hosted on volumes other than those hosting database data and index information, are critical to database operations and must be replicated. In addition, directories holding system files and patches to the database binary files and associated utilities must be replicated to ensure their availability at the remote location.

Summary

By effectively planning and tuning the database server and replication configuration, the amount of data being replicated has been reduced from 250 GB per day to 60 GB per day. In addition, several volumes were identified as critical to operation of the database but unnecessary for replication. This example should be consistent with environments in many IT organizations.

Disaster recovery is not just about data, and it is worth remembering that simply replicating volumes to a DR site will not guarantee the successful recovery of business applications following a disaster event. It is essential to review and analyze the data and files required to return the system to operational functionality. Frequently, as is demonstrated in the scenario detailed here, the demands on data replication can be drastically reduced by applying common sense to the process.

Appendix

VERITA Volume Replicator Advisor (VRadvisor)

VERITAS Volume Replicator Advisor

Additional Parameters for Analysis
Specify the details required for analysis.

Network Bandwidth Available for Replication: Mbps

Secondary Backup Interval: Day(s)

Outage Duration

Network Outage Duration: Hour(s)

Secondary Outage Duration: Hour(s)

Total Outage Duration:

< Back Analyze Cancel Help

Figure 4. Planning the size of the SRL with VRadvisor.

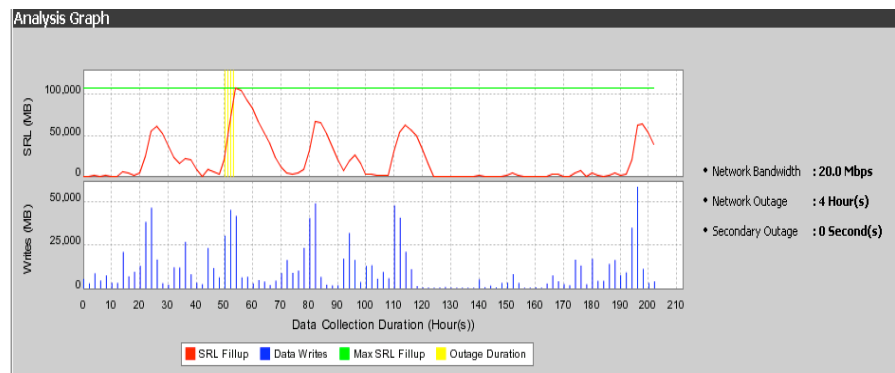


Figure 5. VRadvisor analysis of bandwidth utilization requirements.

VERITAS Software Corporation

Corporate Headquarters
350 Ellis Street
Mountain View, CA 94043
650-527-8000 or 866-837-4827

For additional information about VERITAS Software, its products, or the location of an office near you, please call our corporate headquarters or visit our Web site at www.veritas.com.