Oracle Active Data Guard Far Sync with Axxana Phoenix for Oracle

ORACLE WHITE PAPER | SEPTEMBER 2015

Table of Contents

Executive Summary

Recovering all data after a disaster is the holy grail of disaster recovery. However, very few solutions can reliably deliver on this promise, or do so at a reasonable cost and without negatively impacting database performance.

The Axxana Phoenix System for Oracle Database can be used with Oracle Active Data Guard Far Sync to provide a unique solution for business continuity. When used together they provide both high availability and complete protection for every database transaction at a remote site. They accomplish this in a cost-effective manner and with high performance, regardless of the nature of the disaster or the distance between the sites.

Note: The Axxana Phoenix System for Oracle Database is sold and supported by Axxana. Axxana is solely responsible for all information pertaining to the Axxana Phoenix System for Oracle Database described in this whitepaper. For more information on Axxana Phoenix System for Oracle Database see www.axxana.com.

The Challenges of Obtaining Zero Data Loss Recovery

"Zero Data Loss (ZDL) recovery," means that when a disaster strikes and a failover to a secondary data center takes place, all transactions that were committed at the primary site prior to the disaster can also be recovered at the secondary site. Oracle Data Guard and Active Data Guard in Maximum Availability protection mode ensure that when the primary site fails, a synchronized standby database at the secondary site contains all database transactions that have been committed at the primary site. Data Guard failover to an already running and synchronized copy of the production database quickly restores service and does so with zero data loss.

ZDL using Data Guard and Active Data Guard, however can present the following challenges:

- There can be a negative impact on production performance if the round-trip network latency (network RTT) between sites is too great. This fact is true for any synchronous method of mirroring or replicating data between sites. Network RTT is affected by the distance between the sites, the number of devices in the network path, and by the quality of the communication lines. As network RTT increases, so does the impact on production database throughput and response time. Production database performance in zero data loss configurations can also be impacted if there is insufficient bandwidth available to transmit the required volume of data.
- Additional measures designed to reduce the performance impact of zero data loss protection across long distances can be cost prohibitive. For example, one strategy is to implement a second Data Guard standby database in a separate fully equipped data center, typically within the same metro area as primary site (network RTT of 1-5ms). This enables zero data loss failover to the local copy while the original standby in the remote site continues receives changes asynchronously to provide out-of-region disaster recovery. This addresses many failure scenarios but not all. The close proximity of the synchronous copy will not

address a disaster scenario in which both data centers will experience the same outage. A direct failover to the remote standby database in this scenario will result in data loss.

- Many disaster scenarios involve rolling failures that begin with the failure of the network link between sites. A Data Guard Maximum Availability configuration is intentionally designed to suspend replication when there is a network outage between sites in order to allow new transactions to commit rather than impact the availability of the primary database. Data loss will occur if the primary database were to experience an outage after operating for a period of time in a disconnected mode.

***Note**: There is another Data Guard mode - Maximum Protection, that guarantees all database transactions are protected by a standby even in the case of multiple failures (e.g. first a network outage, and then a primary database outage). However this also requires additional investment to deploy two complete synchronous standby databases in order to prevent the inability to communicate with one of the standby's from affecting the availability of the production database.*

The challenges of implementing zero data loss recovery frequently cause customers to compromise on data protection. Rather than incur the impact to database performance or the additional cost and complexity of various strategies to mitigate the impact, they simply use Data Guard Maximum Performance mode with asynchronous replication. The upside of asynchronous replication is that it has no impact on database response time or throughput because there is never any wait for confirmation from a standby database that data is protected. This performance advantage comes at a cost, however, because asynchronous replication is unable to provide a zero data loss guarantee. Zero data loss failover has two key advantages, one of which applies to just about every customer having a mission critical database:

- There are many cases where transactional data is so critical to an enterprise or its customers that it is undesirable to lose even a single transaction. Zero data loss protection is a fundamental requirement for these enterprises.
- There are many other cases where an enterprise can tolerate a limited amount of data loss. Even then, however, there is usually a preference not to lose data or at a minimum to understand what data was lost. This often results in behavior that contradicts best practices for high availability. For less catastrophic outages when it is possible to eventually repair and restore a database, businesses will often postpone failover to their business continuity site until they take every possible measure to recover data. This will extend the length of an outage in a very unpredictable way even though they could have failed over immediately and quickly restored service. And then there are truly catastrophic outages that result in the primary database being unrecoverable. With no ability to recover the primary database it is impossible to know which transactions were not replicated before the outage occurred.

For these reasons Oracle introduced Active Data Guard Far Sync with Oracle Database 12c to cost-effectively address the requirement of zero data loss failover to a synchronized standby database deployed at a remote location, whether it be located 100's or 1000's of miles away from the production database.

Oracle Active Data Guard Far Sync

Active Data Guard Far Sync is a new capability with Oracle Database 12c that eliminates compromise by enabling zero data loss failover to a standby database located at any distance from the primary site. Far Sync does this without impacting the performance of the production database or requiring multiple replicas that increase cost and operational complexity.

Far Sync enables zero data loss failover at any distance by deploying a Far Sync instance (a lightweight Oracle instance that has only a control file, spfile, password file and standby log files, there are no database files or online redo logs) at a distance that is within an acceptable range of the primary for synchronous transport (SYNC). A Far Sync instance receives redo from the primary database synchronously and immediately forwards it to up to 29 remote standby databases using asynchronous transport (ASYNC) as described in Figure 1.

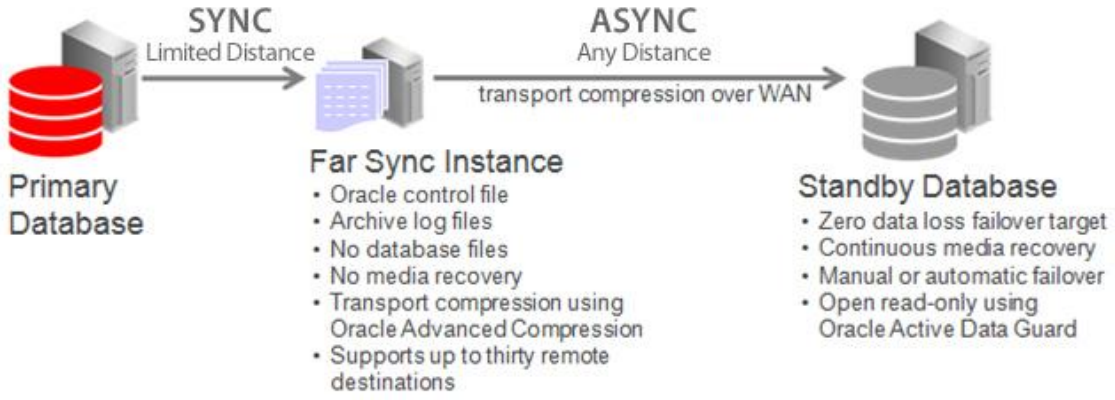

Figure 1: Active Data Guard Far Sync Architecture

The presence of a Far Sync instance in a Data Guard configuration is transparent to its operation during switchover or failover, the administrator uses the same commands used for any Data Guard configuration to perform a zero data loss failover to the remote standby database. Far Sync requires nothing new to learn or any additional procedures in order to perform a zero data loss failover across a wide area network (WAN).

In addition to enabling zero data loss failover at any distance, a Far Sync instance offloads the primary database of any overhead for the following tasks:

- Resolving gaps in archived logs received by the remote standby database(s) (e.g. following network or standby database outages).
- Redo transport overhead for configurations having multiple standby databases. The primary ships once to the Far Sync instance, Far Sync takes care of shipping to multiple destinations.
- A Far Sync instance can also perform off-host network compression, conserving WAN bandwidth without impacting primary database performance.

The lightweight nature and transparent operation of the Far Sync instance is a substantial improvement over the previous multi-standby solution to WAN zero data loss protection; there are no user data files, no media recovery and no Oracle Database license required for the Far Sync instance. The Far Sync instance only needs sufficient disk space for standby redo log files and to retain archived redo logs to resolve any gaps that might occur. The Far Sync instance requires a very small SGA footprint (much less than production) and consumes less than one CPU (except when it also performs redo transport compression).

Far Sync provides increased flexibility in the location of a disaster recovery site for those who wish to implement zero data loss protection. Even users who have already deployed Data Guard synchronous transport can benefit from configuring a Far Sync instance closer to the primary than their current standby to reduce the performance impact on the production database.

Far Sync also benefits users who currently use asynchronous transport. Upgrading to zero data loss protection using Far Sync eliminates the uncertainty and administrative burden that accompanies the need to reconcile data loss after a failover has occurred in an asynchronous configuration. Far Sync can increase availability by eliminating the tendency to postpone failover in the hope of resolving an outage without losing data. Far Sync's guarantee of zero data loss encourages immediate failover to quickly resume service while the problems that caused the outage are resolved. The primary database in an ASYNC configuration also enjoys the same benefits of offloading the overhead of gap resolution, servicing multiple standby databases, and redo transport compression.

In almost every case, Far Sync is ideal for achieving the highest level of data protection with the least performance impact while also lowering network bandwidth consumption.

Far Sync at its core is a clever forwarding mechanism to enable zero data loss failover at any distance – integrated, lightweight, low cost, and simple to operate. What Far Sync can't do is defy physics. Synchronous replication between the production database and a Far Sync instance mean they have to be located within the same metro region even if the standby database can now be located on the other side of the world. A disaster that affects the metro area causing an outage of both the production database and the Far Sync instance, or rolling disasters that first impact the network then impacts the production system can still result in data loss. Deploying a Far Sync instance in a separate location from the production database can also impact cost. Addressing these remaining challenges is what makes the Axxana Phoenix System an ideal complement to Active Data Guard Far Sync.

Axxana Phoenix System for Oracle

Axxana's Phoenix System is the first true disaster-proof storage and server unit bundled together. It is composed of a disaster-proof storage system and a Linux-based server enclosed in a hardened case (Axxana's Black Box) that is capable of surviving extreme conditions, such as:

- Direct flames up to 20000 F
- Enormous pressures associated with building collapse
- Pierce forces associated with falling rods (specifically, a pierce force of a 500 lb. rod with a cross-section of 0.04 in2 dropped from a height of 10 ft.)
- Long periods of submersion underwater
- Extreme shocks

In short, data stored in the Phoenix System with its Axxana Black Box will survive any disaster short of an atomic explosion.

In addition to its sturdiness and survival features, the Axxana Black Box has built-in independent capabilities to quickly transfer data out of the Black Box.

The Black Box's data transfer capabilities include:

- Self-sufficient power, as most disaster scenarios will result in the loss of the external power supply. The Black Box is equipped with an internal battery, enabling it to continue operation for up to 36 hours post disaster. This is needed to enable the Black Box to communicate with the secondary data center and transfer the protected data.
- A Linux-based server.
- Broadband LTE transmission capability to transmit the data over an LTE cellular network.

- Wi-Fi transmission capability to facilitate data extraction to a nearby laptop in case the cellular transmission fails.

Figure 2 depicts The Phoenix System with its Axxana Black Box, showing the complete system that encases the data-protection chamber. Through the transparent side of the enclosure one can see the fire-protection box. The fire-protection box combines multiple materials, technologies, and techniques that enable the device to withstand extreme heat and direct flames, and at the same time permit cellular transmission of data. The data-protection chamber houses the protected server and the storage. The chamber is made of thick steel to provide protection from mechanical threats. In this picture one can also see two of the antennas, which are themselves designed to withstand high temperatures. The chamber is cooled by the pipes encircling the cylinder. Cooling is necessary to dissipate the heat generated by the internal electronics under normal operating conditions. The internal electronics within the chamber consist of solid-state drives capable of withstanding high levels of shock. These drives are designed to resist mechanical shocks far more resiliently than rotating drives.

Figure 2: Axxana Phoenix System External View with the Axxana Black Box Internals

In summary, The Phoenix System provides a new concept in data protection. In a sense, it is a bunker in a box with independent communication capabilities built to survive disasters.

Active Data Guard Far Sync with Phoenix System for Oracle Solution

Active Data Guard Far Sync is a flavor of multi-site replication in which the active database at the primary site synchronously replicates to the Far Sync instance at the nearby site, which in turn replicates asynchronously to a remote standby database.

Although the Far Sync instance is a very low-footprint, low-cost instance, the solution has to be bound by the limitations described in the previous section: the sensitivity to communication failures, latency, performance impact, cost of the nearby site, and vulnerability to more-global disasters are all still there. However, when combined with Axxana Phoenix System, Far Sync is able to overcome all of these limitations.

Figure 3 illustrates the use of Axxana Phoenix System with Far Sync. The Far Sync instance is deployed on the Linux system and storage embedded in the Axxana Black Box just as it would be on any other server. While Axxana provides a completely different level of protection, the operation of the Far Sync instance is the same as on any other server.

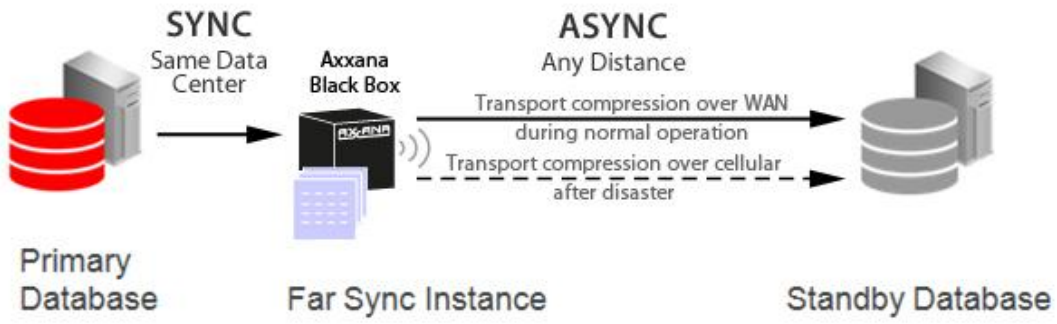

Figure 3: Active Data Guard Far Sync Architecture with Axxana Phoenix System for Oracle and its Black Box

Although Figure 3 shows the Axxanna system hosting a Far Sync instance for a single Active Data Guard configuration, it can also house Far Sync Instances for other Active Data Guard configurations, providing zero data loss protection for multiple production databases.

If there is a production database outage, failover to the remote standby database can happen automatically using Data Guard Fast-Start Failover or can be manually initiated by the administrator – the same as when Far Sync is deployed on a generic server. The Axxana system will even re-establish a Data Guard session with the standby instance at the remote site over an independent cellular communication line if necessary. Active Data Guard will complete the data transfer of all remaining redo records from the Far Sync instance , apply them to the standby database, and reconfigure it as the new primary database. The new primary database will have all transactions committed to the original primary database prior to the disaster resulting in zero data loss.

In an extreme event where even the cellular network is down, another recovery option is to retrieve the information stored in the Axxana Black Box to a laptop using Wi-Fi and running the Far Sync instance on the laptop, which can then be moved to a location where Data Guard connectivity with the standby database can be established before initiating the failover process.

Since the Axxana Phoenix System is immune to disaster scenarios, it can reside at the primary site right next to the primary database servers. Thus there are no additional costs associated with a nearby site and no unnecessary network latency that can negatively impact production database performance.

The Axxana Phoenix System increases the probability of ZDL with Oracle Far Sync solution for two reasons:

- The Far Sync instance will survive regional disasters by running on the Axxana system. This is not possible in cases where a Far Sync instance is deployed on a generic server within the same region as the primary (required for synchronous communication). A regional disaster is likely to affect the availability of both the primary system and the Far Sync instance and result in data loss. Data loss is prevented when

the Far Sync instance is deployed on Axxanna because the Axxanna system will survive the disaster and still be able to transmit data.

- Far Sync in the Axxana Phoenix System solution is much more resistant to communication failures than Far Sync deployed on a generic server. This is true on two counts:
    - The Far Sync instance is now located at the primary site, so there are no communication issues to the nearby separate site as previously described.
    - If communication to the standby site fails first, Axxana Phoenix System has sufficient storage to continue to collect and protect Oracle redo. Depending on the redo generation rate of the primary database this can continue for a number of minutes (potentially 15 to 30 minutes), beyond the time necessary to determine whether a disaster is in progress.

Conclusion

The immunity of the Axxana Phoenix System to most disaster scenarios combined with Oracle Active Data Guard Far Sync provides the most comprehensive, reliable, and cost-effective zero data loss recovery and high availability solution for enterprise-critical data.

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

ORACLE

CONNECT WITH US

blogs.oracle.com/oracle
facebook.com/oracle
twitter.com/oracle
oracle.com

Integrated Cloud Applications & Platform Services

Oracle Active Data Guard Far Sync with Axxana Phoenix for Oracle
September 2015
Author: Larry M. Carpenter
Contributing Authors: Alex Winokur, Axxana