AXXANA

BUILT TO LAST

Squeezing the Max out of 12c Disaster Recovery

Alain Azagury, VP R&D
Yossi Nixon, Oracle ACE Associate & Chief DBA
November2017

ORACLE Gold Partner Cloud Standard

ORACLE DATABASE READY

ORACLE LINUX READY

ORACLE BIG DATA APPLIANCE READY

ORACLE EXADATA READY

© All Rights Reserved. Axxana 2017

Speakers

Alain Azagury, VP R&D

- 30-years of experience in Storage, Virtualization, Cloud, Systems, Memory Management, etc.
- IBM Master Inventor, with over 15 patents
- Avid cyclist

Yossi Nixon, Chief DBA

- Oracle ACE Associate, Database Architect and Oracle DBA with 20 years of experience.
  Specialties: DBA & Infrastructure Architect; Disaster Recovery Strategy Development; Monitoring and troubleshooting; Designing and Writing code....

Confidential

© All Rights Reserved. Axxana 2017

Agenda

- Data Guard Overview
- Data Guard Terms
- Active Data Guard Far Sync
- Data Guard DR enhancements in Oracle 12c
- Lessons Learned
- The Risk-Distance Conundrum
- Axxana's Phoenix Technology
- Putting it all together

© All Rights Reserved. Axxana 2007-2017
CONFIDENTIAL

© All Rights Reserved. Axxana 2017

Data Guard Overview

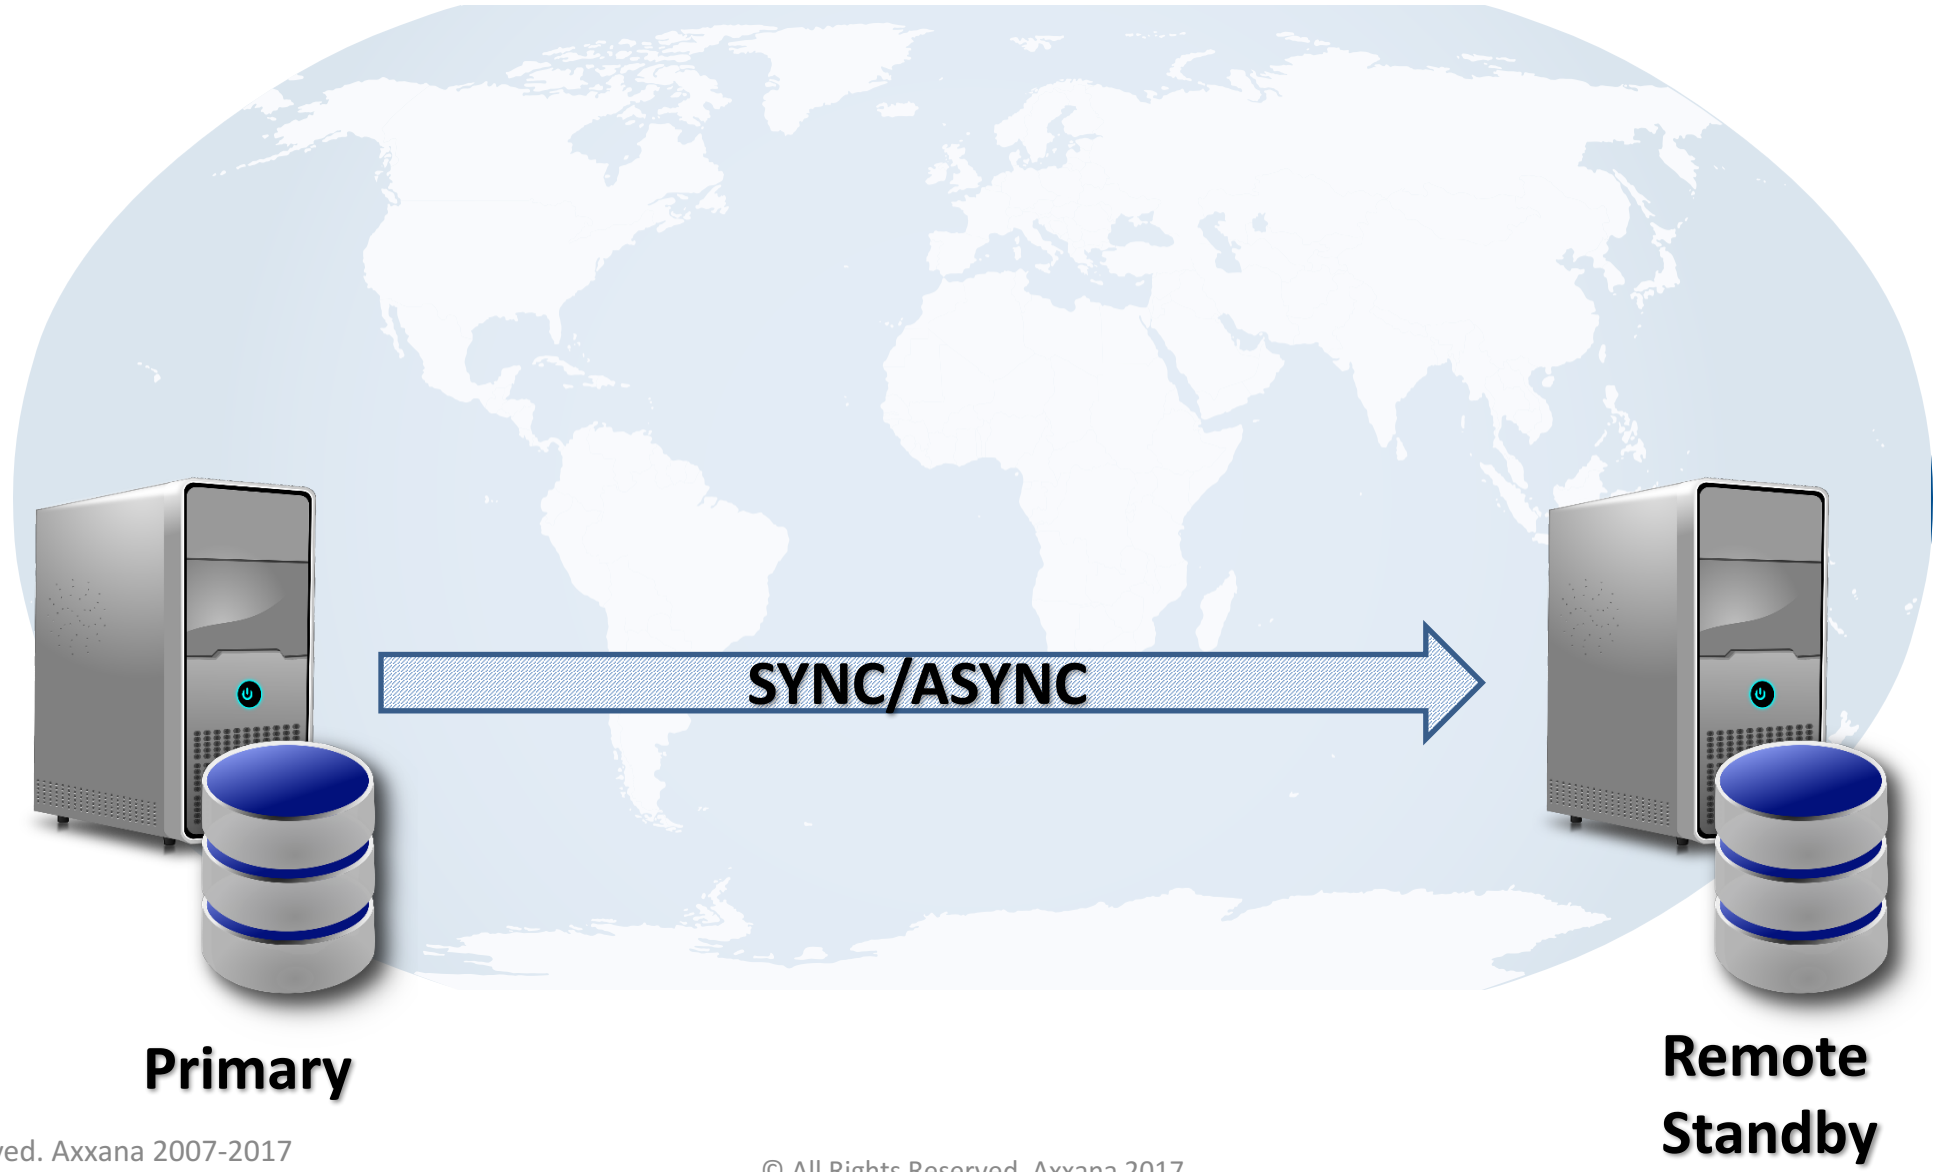

© All Rights Reserved. Axxana 2007-2017
CONFIDENTIAL

© All Rights Reserved. Axxana 2017

Data Guard Terms

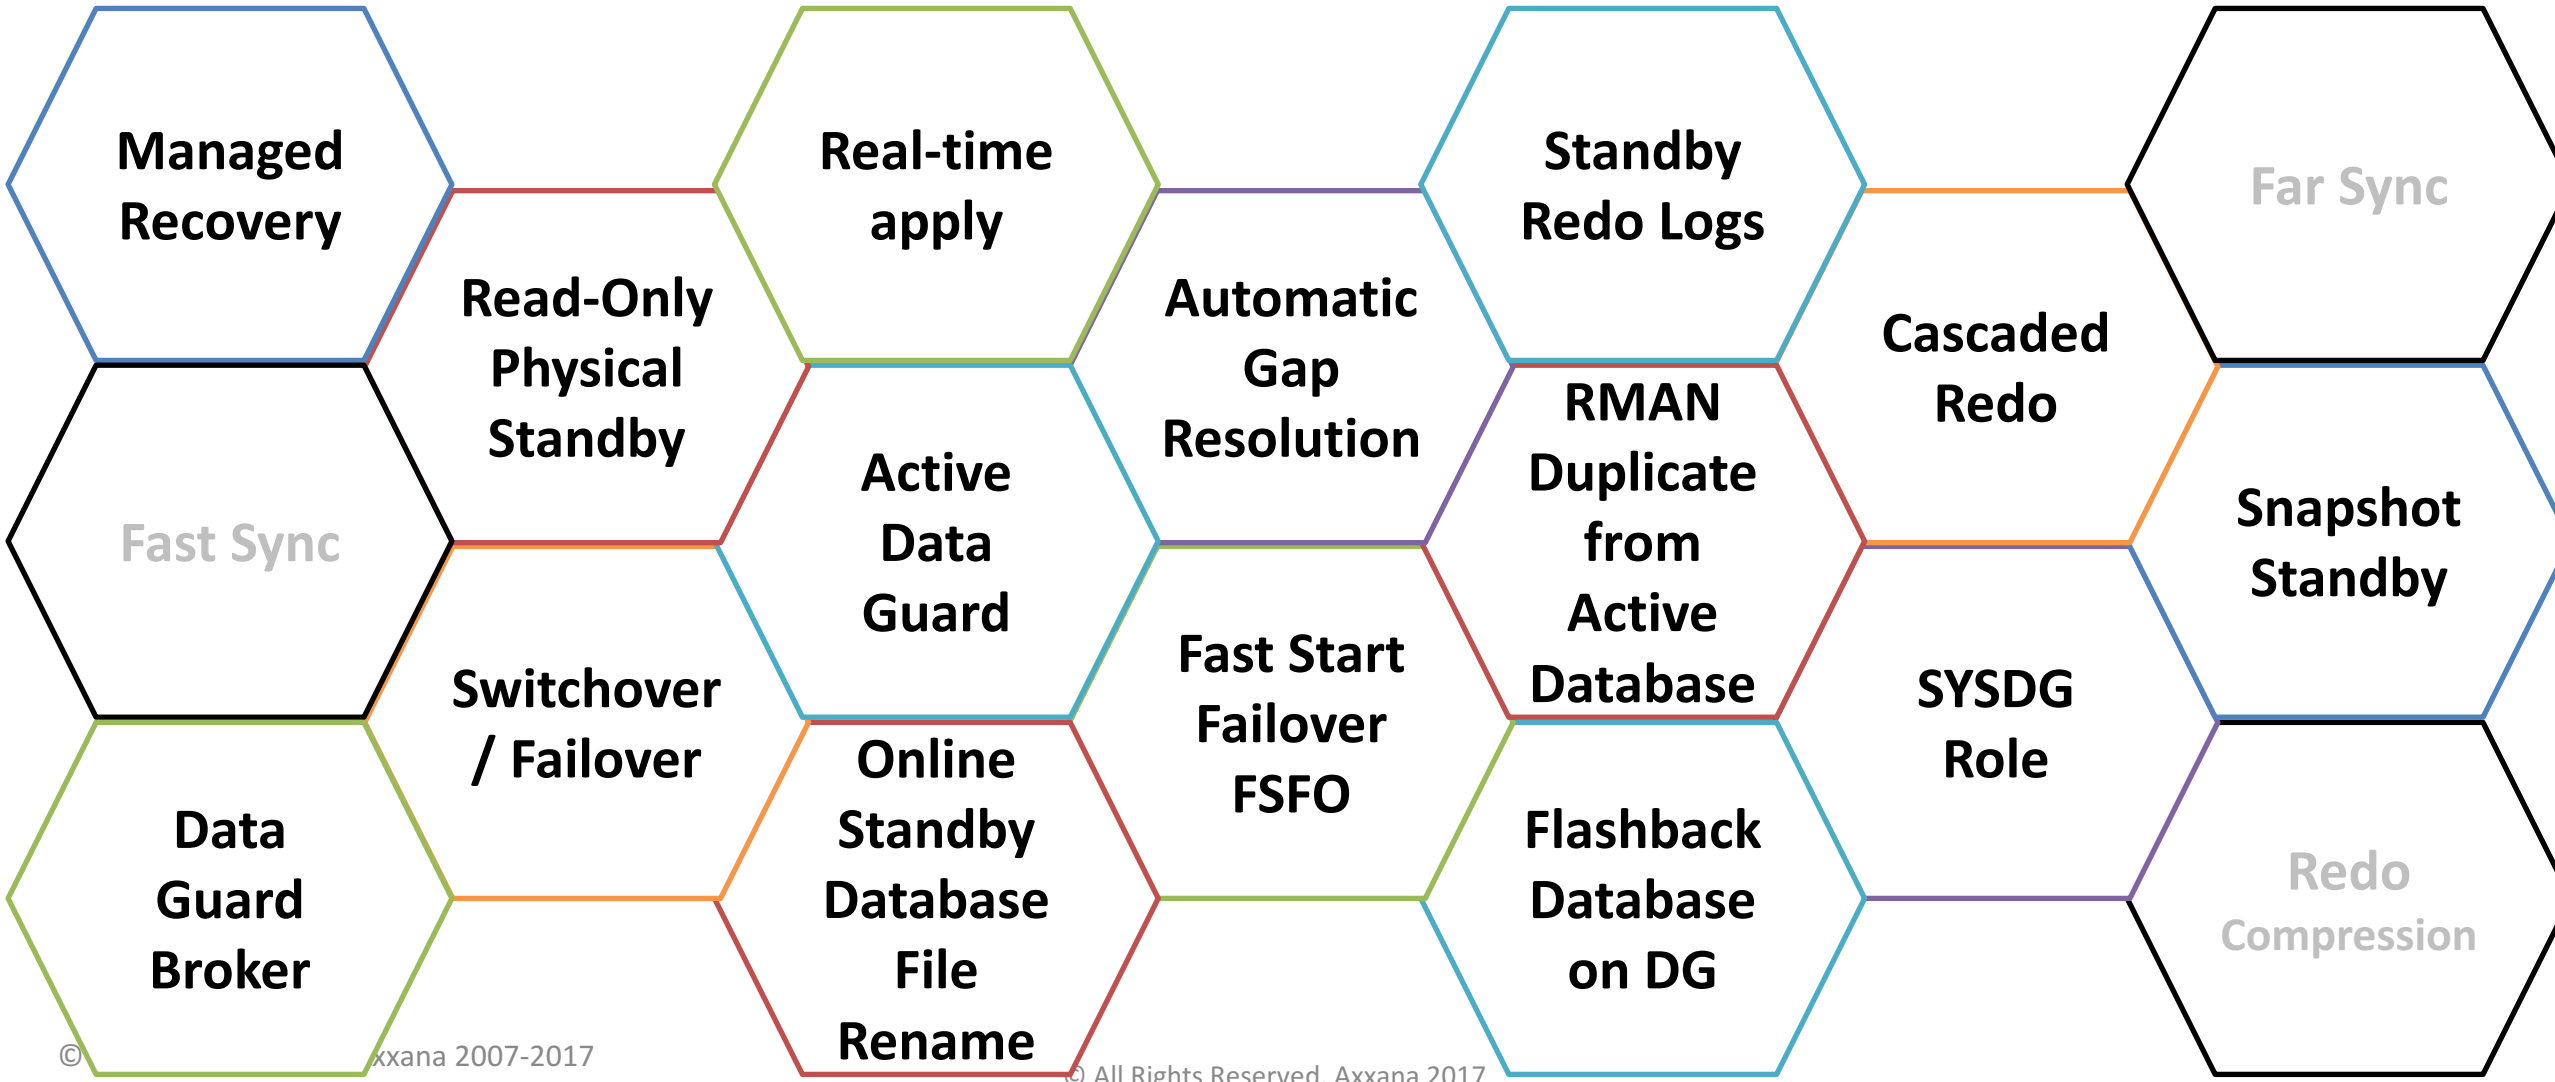

© xxana 2007-2017
CONFIDENTIAL
© All Rights Reserved. Axxana 2017

Active Data Guard Far Sync

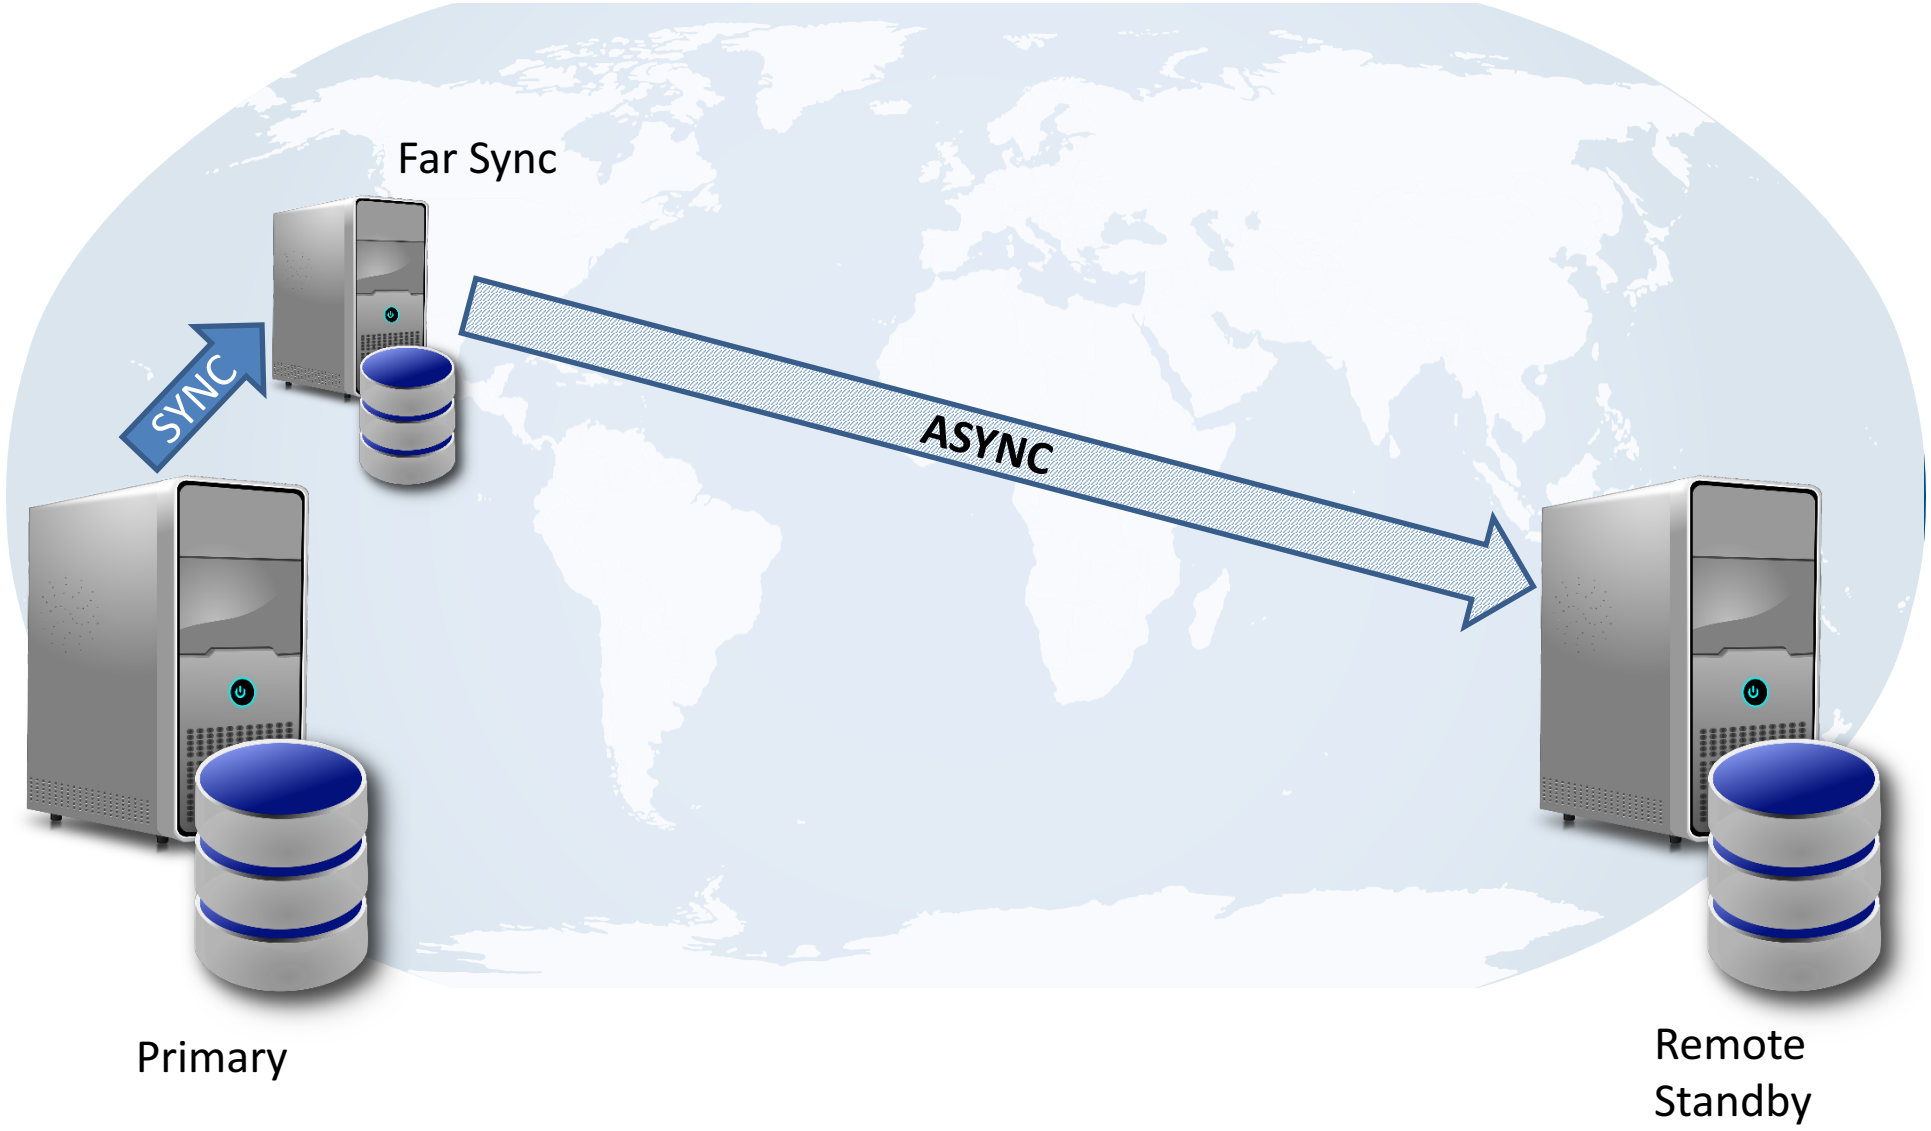

Confidential

© All Rights Reserved. Axxana 2017

Active Data Guard Far Sync

- Far Sync: light-weight Oracle instance: standby control file, standby redo logs, archived redo logs, no data files
- Receives redo synchronously from primary, forwards redo asynchronously in real-time to standby
- Upon Failover: Async standby transparently obtains last committed redo from Far Sync and applies: zero data loss failover
- Minimal CPU, memory, or I/O– No recovery, no data files, production never runs on the Far Sync instance

Confidential

© All Rights Reserved. Axxana 2017

Far Sync summary

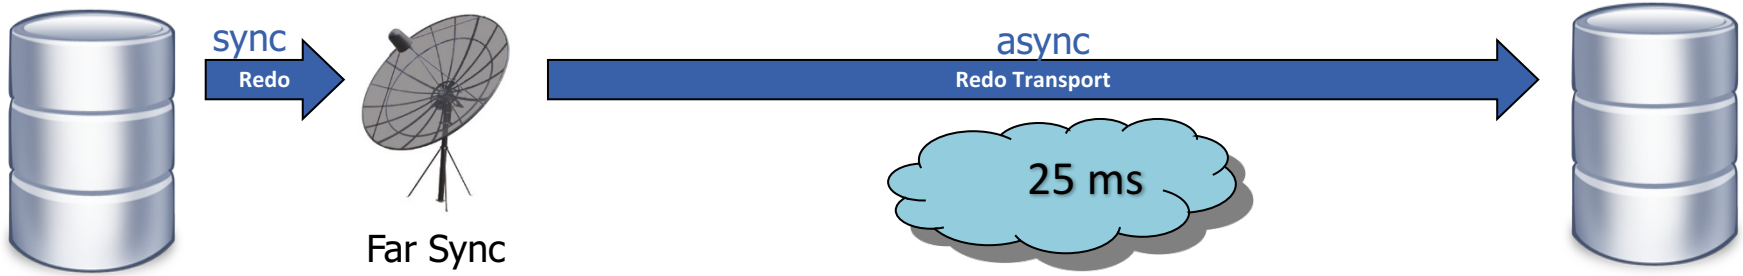

- Very promising for long distance replication if data loss is not acceptable
- Up to 60% performance gain (DML only workloads) with 25ms network latency
- Lightweight and easy to deploy (virtual machine)
- If latency <5ms most likely you don't need Far Sync
- There are still bugs that need fixing

© All Rights Reserved. Axxana 2007-2017
CONFIDENTIAL
© All Rights Reserved. Axxana 2017

Far Sync Syntax

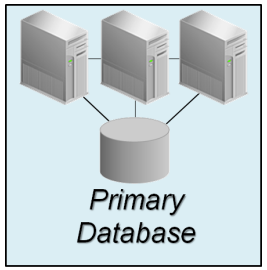


```
SQL > ALTER DATABASE CREATE FAR SYNC INSTANCE
CONTROLFILE AS  '/tmp/farsync.ctl';
SQL > create pfile='/tmp/initfs.ora' from spfile;
```

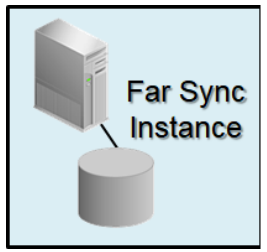


```
SQL > ALTER DATABASE ADD STANDBY LOGFILE THREAD 2
SIZE 52428800;
SQL > create spfile from pfile='/tmp/initfs.ora';
```

```
DGMGRL> ADD FAR_SYNC fs AS CONNECT IDENTIFIER IS fs;
```

```
DGMGRL> EDIT DATABASE pdb SET PROPERTY RedoRoutes = '(LOCAL
: fs SYNC ALT =(sdb ASYNC FALLBACK))';
```

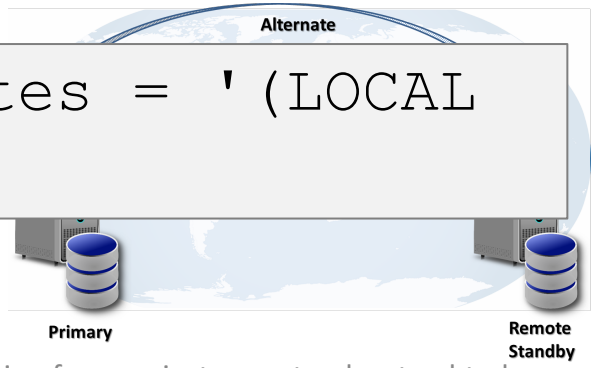

© All Rights Reserved. Axxana 2007-2017
CONFIDENTIAL

© All Rights Reserved. Axxana 2007-2017
https://ora-10dba.blogspot.com/2017/08/creating-far-sync-instance-step-by-step.html

Lab Environment Setup

- Oracle 12.1 Enterprise Edition
  - Active Dataguard with Far Sync
  - ASM
  - Requested tailored patches for Far Sync
  - Oracle Linux 7.2
- Server environment
  - Primary – 2 x Lenovo X3550 M5 (2 x Xeon E5-2620 V3 2.4Ghz, 32GB memory, 4 SSDs)
  - Secondary – Identical to Primary
  - Far Sync – Two configurations
    - i7/16GB/dual SSD or,
    - Quad Xeon/32GB/4 SSDs
- Neutralized network latency issues
  - 10GbE connectivity
  - All systems are collocated
- 1 to 16 instances
  - Up to 8 instances on each primary/secondary server
  - Up to 16 instances on Far Sync server

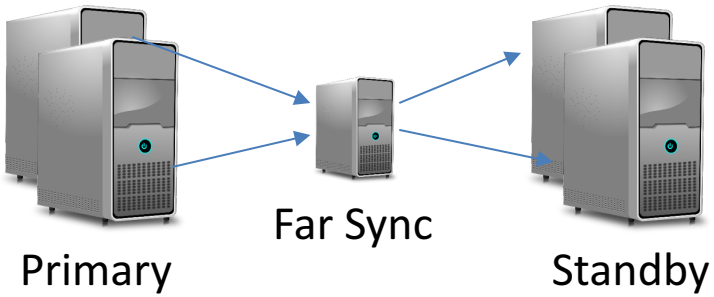

© All Rights Reserved. Axxana 2017

Setup Observations

- Alternate
- Opposite direction
- Maximum Availability
- No multiplex
- RMAN deletion policy

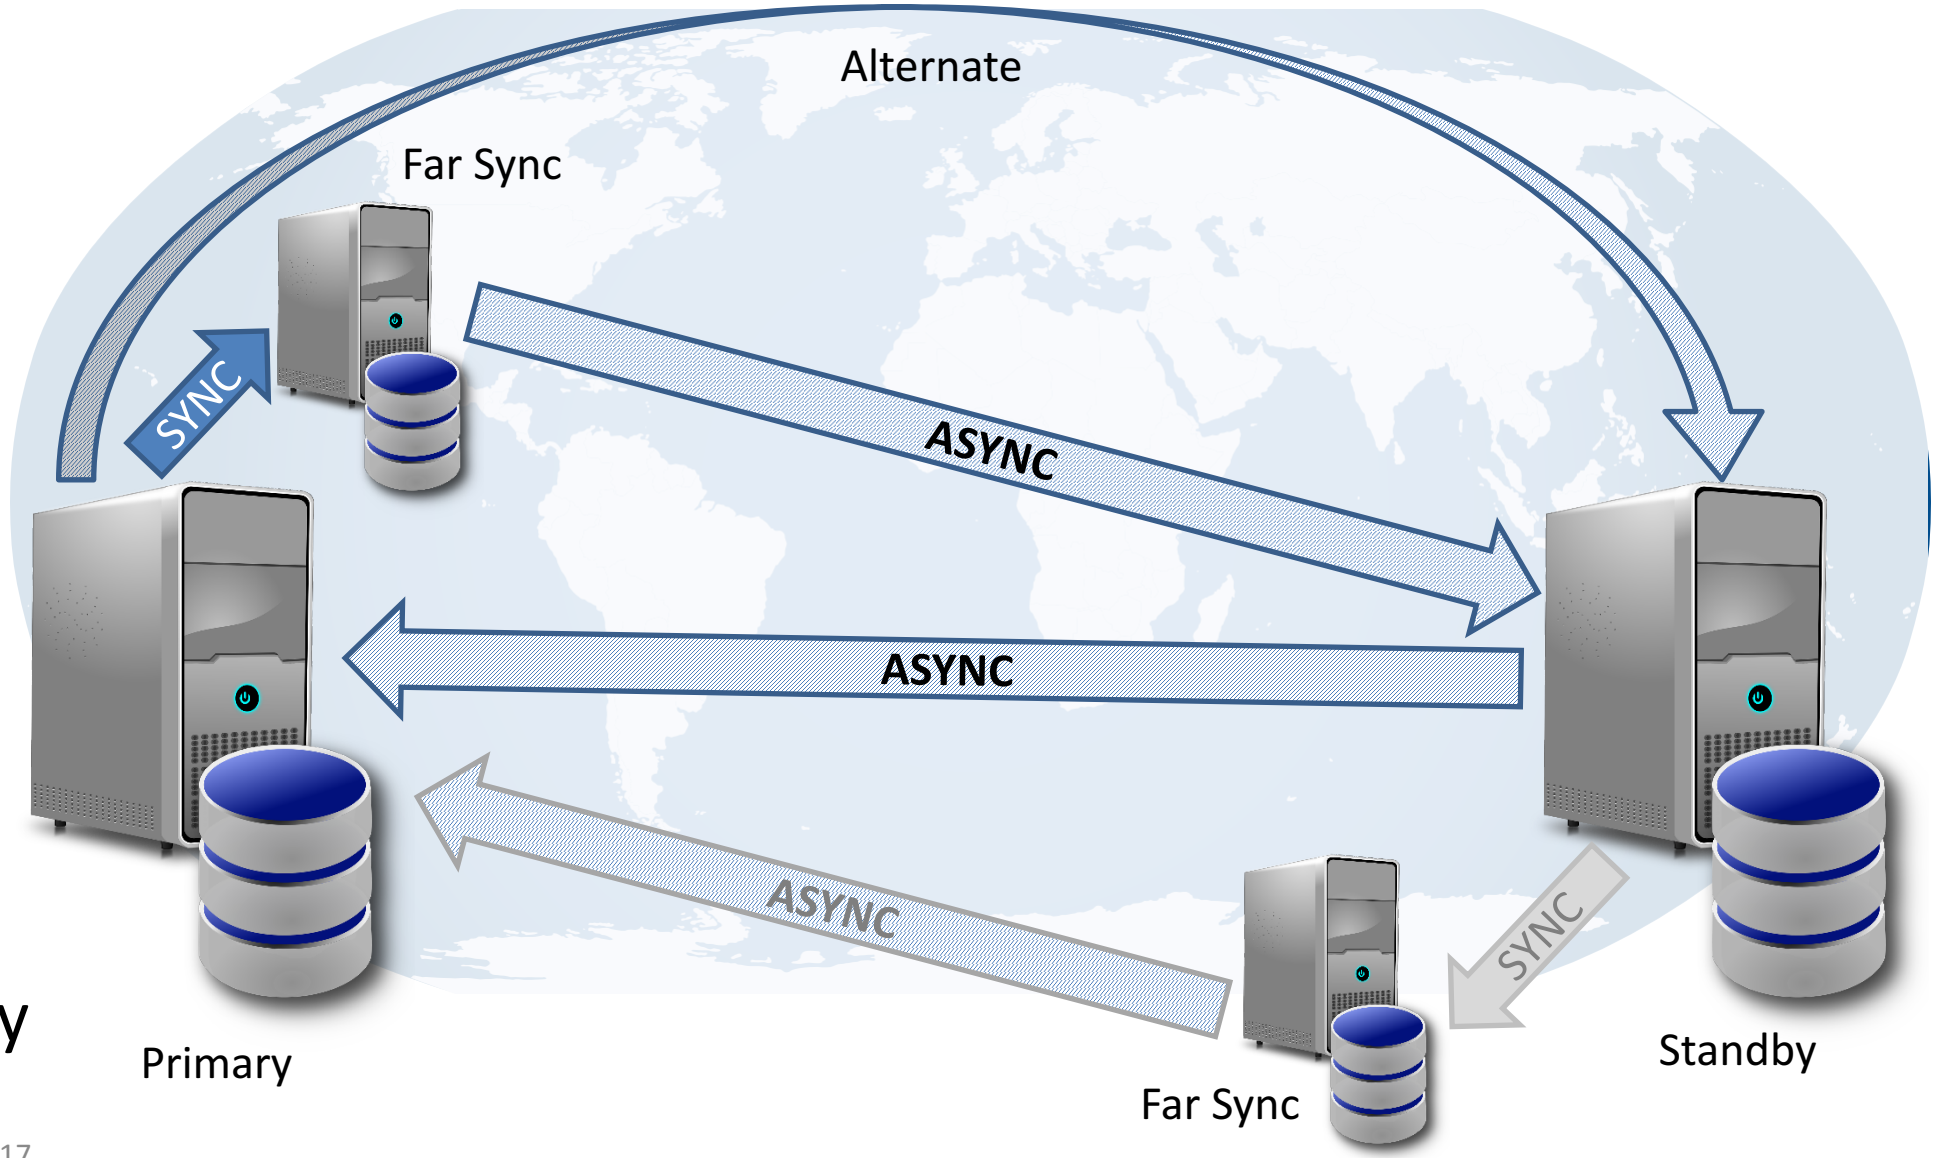

© All Rights Reserved. Axxana 2007-2017
CONFIDENTIAL

© All Rights Reserved. Axxana 2017

Lessons Learned

- ✅ Memory Consumption
- ✅ Networking
- ✅ Latency reduction through Fast Sync
- ✅ CPU Consumption
- ✅ Far Sync IO patterns

Confidential

© All Rights Reserved. Axxana 2017

Memory Consumption

- **Configuration**
  - Used the recommended SGA allocation
    - 300MB SGA
- The Far Sync instance requires a very small SGA footprint (much less than production)
  - No database files
  - No media recovery
- Consider allocating memory for storage cache (later…)

Confidential

© All Rights Reserved. Axxana 2017

Networking

- **Configuration**
  - Jumbo frames make a huge difference
    - We used MTU 9000
  - Bonding two 10GbE ports
- Network Traffic is Consistent with Log File Write throughput
  - Bytes written = Bytes received by Far Sync Instance = Bytes Sent by Far Sync Instance
- Tune LOG_ARCHIVE_MAX_PROCESSES to maximize throughput
- Transport compression will reduce traffic… but takes its toll on CPU cycles…

Confidential

© All Rights Reserved. Axxana 2017

Latency Reduction – Fast Sync

- Standby acknowledges the primary as soon as it receives redo in memory
  - Without waiting for disk I/O to a standby redo log file
- Main goal is shortening the total round-trip time between primary and standby
- We didn't see yet significant benefit from Fast Sync (NOAFFIRM), need more experimentation!

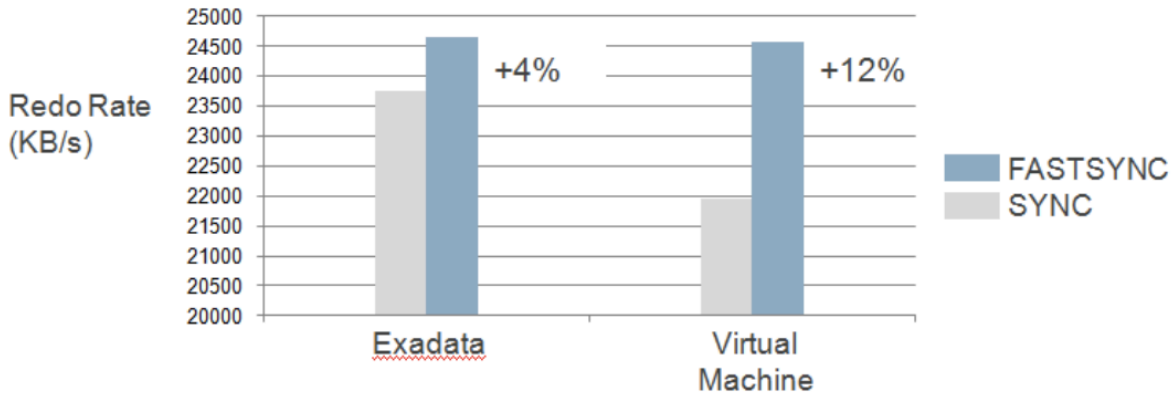

Confidential

CPU Consumption

- Followed Oracle's recommendation
  - Reducing the CPU_COUNT during testing had no effect on the performance of the Far sync instance
- Virtual Keeper of Time Process(es) (VKTM) are high CPU consumer
  - One process per Database and ASM instances
  - Several "solutions" have been recommended to reduce its priority

    ✅**Significant improvement achieved when Far Sync runs multiple instances!**
- Transport compression takes its toll on CPU cycles…

Confidential

© All Rights Reserved. Axxana 2017

Far Sync IO Patterns

- Redo entries stored in Standby Redo Logs and Archives
- Far Sync does not take advantage of log_buffer memory to reduce I/O
  - IOs may be reduced through smart caching techniques…
- Reduce Redo Log members

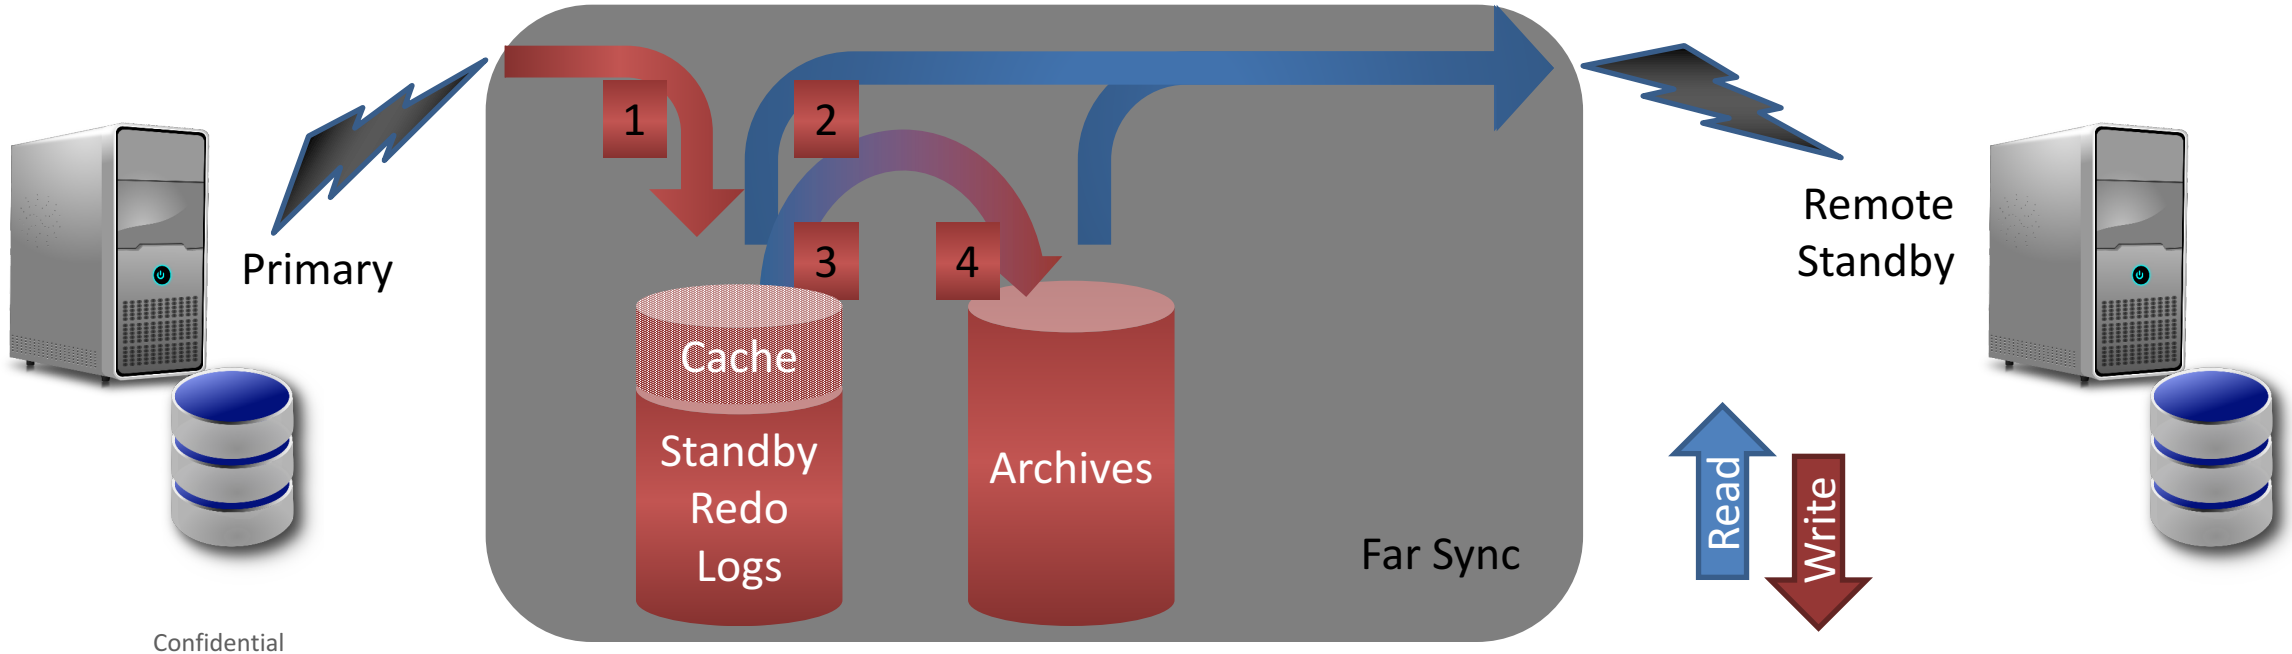

Confidential

Additional Observations

- Alternate
- Opposite direction
- Maximum Availability
- No multiplex
- RMAN deletion policy

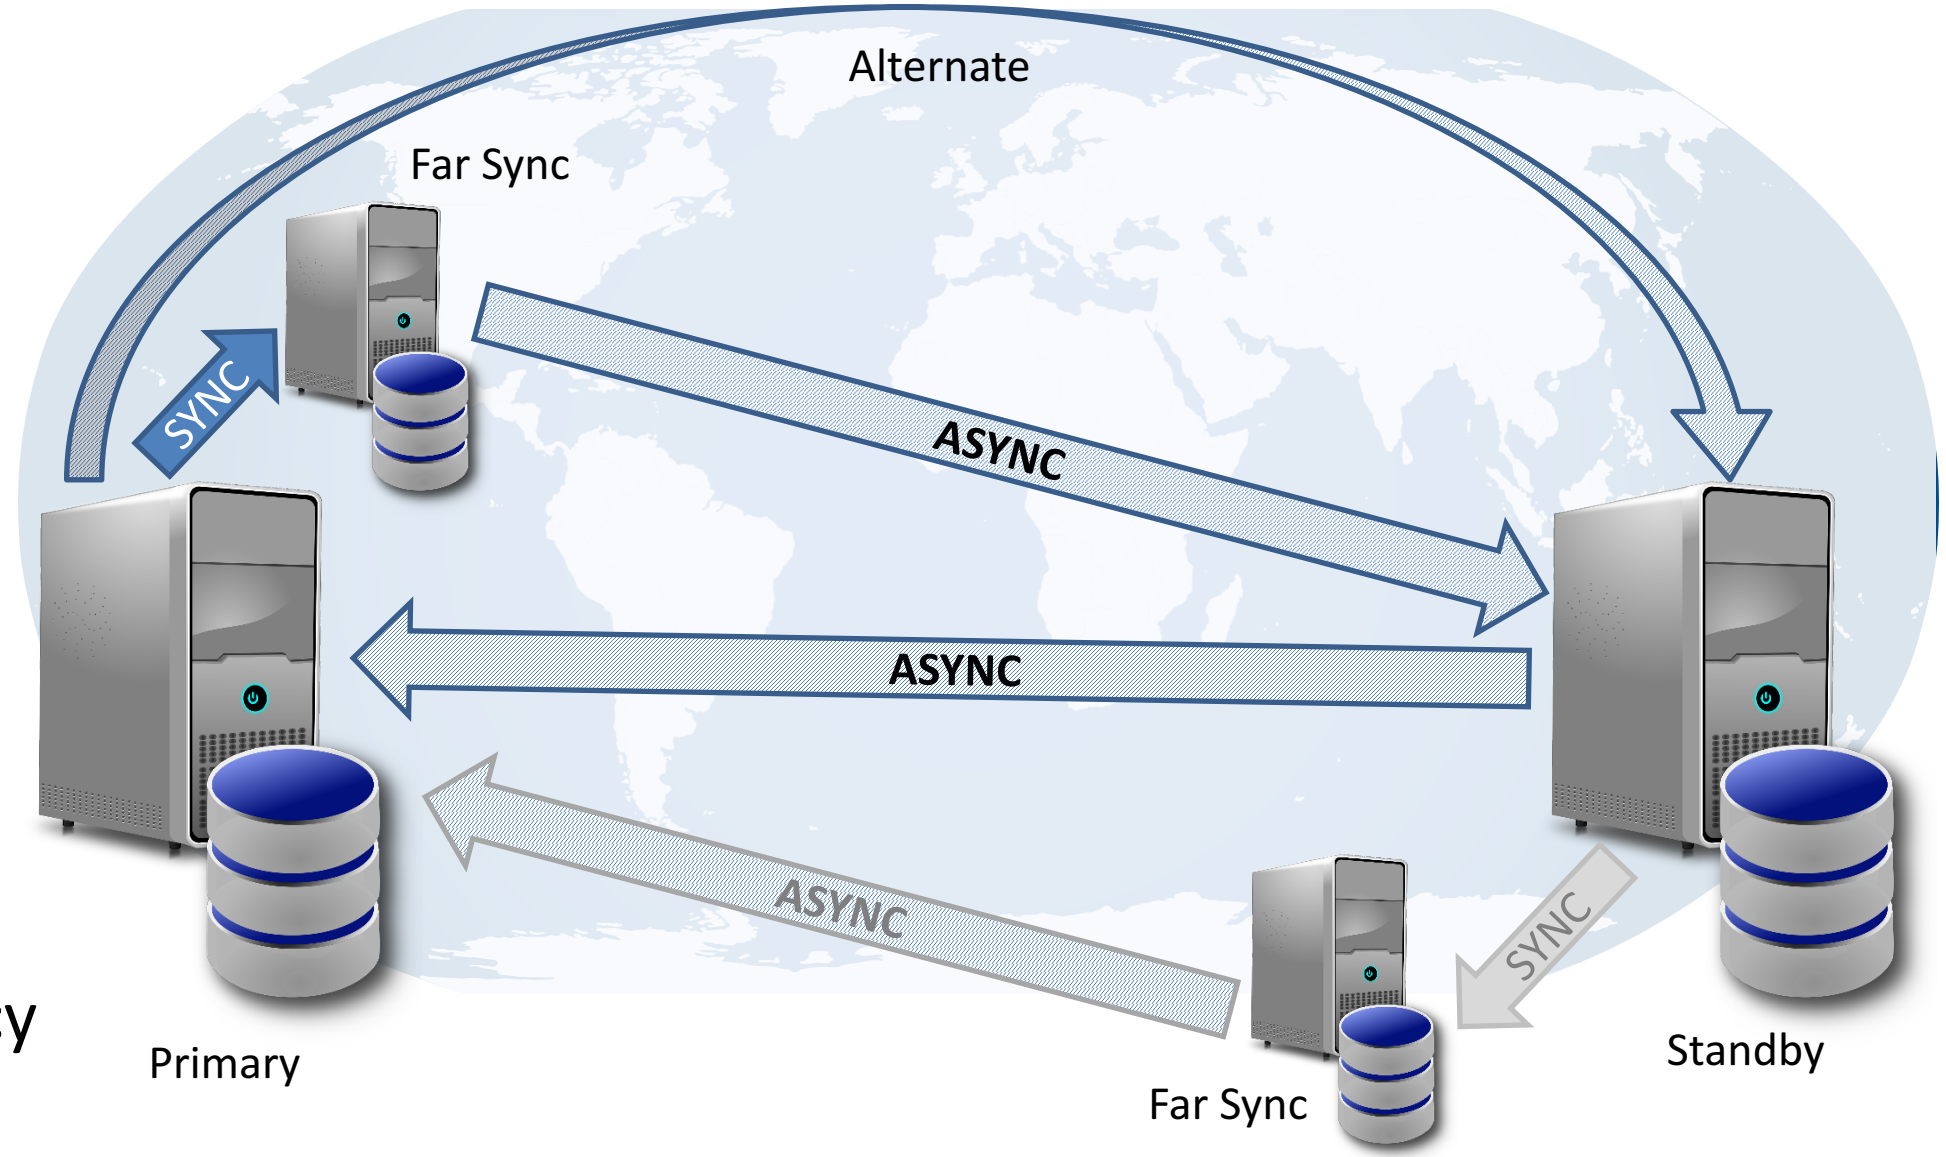

Confidential
© All Rights Reserved. Axxana 2017

Additional Observations *(cont.)*

- Using Data Guard Broker is strongly recommended, especially when using Far Sync
- No “easy” way to collect database statistics, since Far Sync is not a database!
- When using Far Sync, use the latest patchsets
- Rolling Forward a Physical Standby Database Using the RECOVER FROM SERVICE is cool (First introduced in Oracle 12.1)

Confidential

© All Rights Reserved. Axxana 2017

Oracle 12.2 New Features

- Block Comparison tool in DGMGRL
- Multiple Observers
- Supports AWR on Active DG for the standby
- Data Guard Standby Creation with dbca
- Multi-Instance Redo Apply
- Fast Failover - Data Guard takes over session draining
- Password file is managed and transported via the Redo mechanism
- Alternate prioritization

Confidential

© All Rights Reserved. Axxana 2017

Oracle 18c New Features

- In Memory Column store works with multi-instance redo apply
- Multi-instance redo apply support with Block Change Tracking (BCT).
- Global Temporary Table support from Standby.
- DML operations on Standby – Redirection to Primary without ACID compromise.
- Preserve buffer cache during role-change.
- No logging enhancement - Performance mode and Availability mode.
- **RMAN** - Refresh a standby from primary/backup using single RECOVER command

© All Rights Reserved. Axxana 2007-2017
CONFIDENTIAL

© All Rights Reserved. Axxana 2017

Distance-Risk Conundrum

- Distance Between Primary and Far Sync

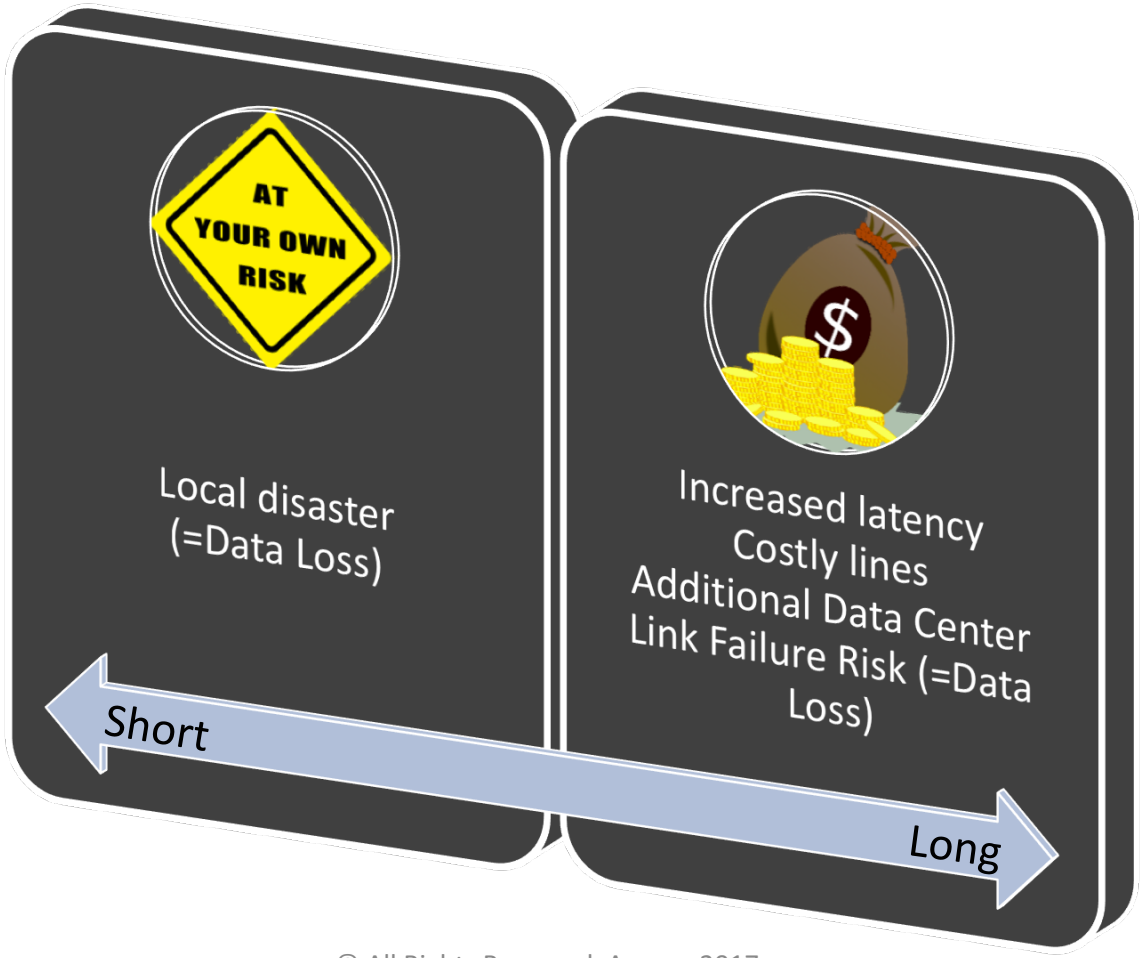

Confidential

© All Rights Reserved. Axxana 2017

Axxana – Transparent Layered Protection

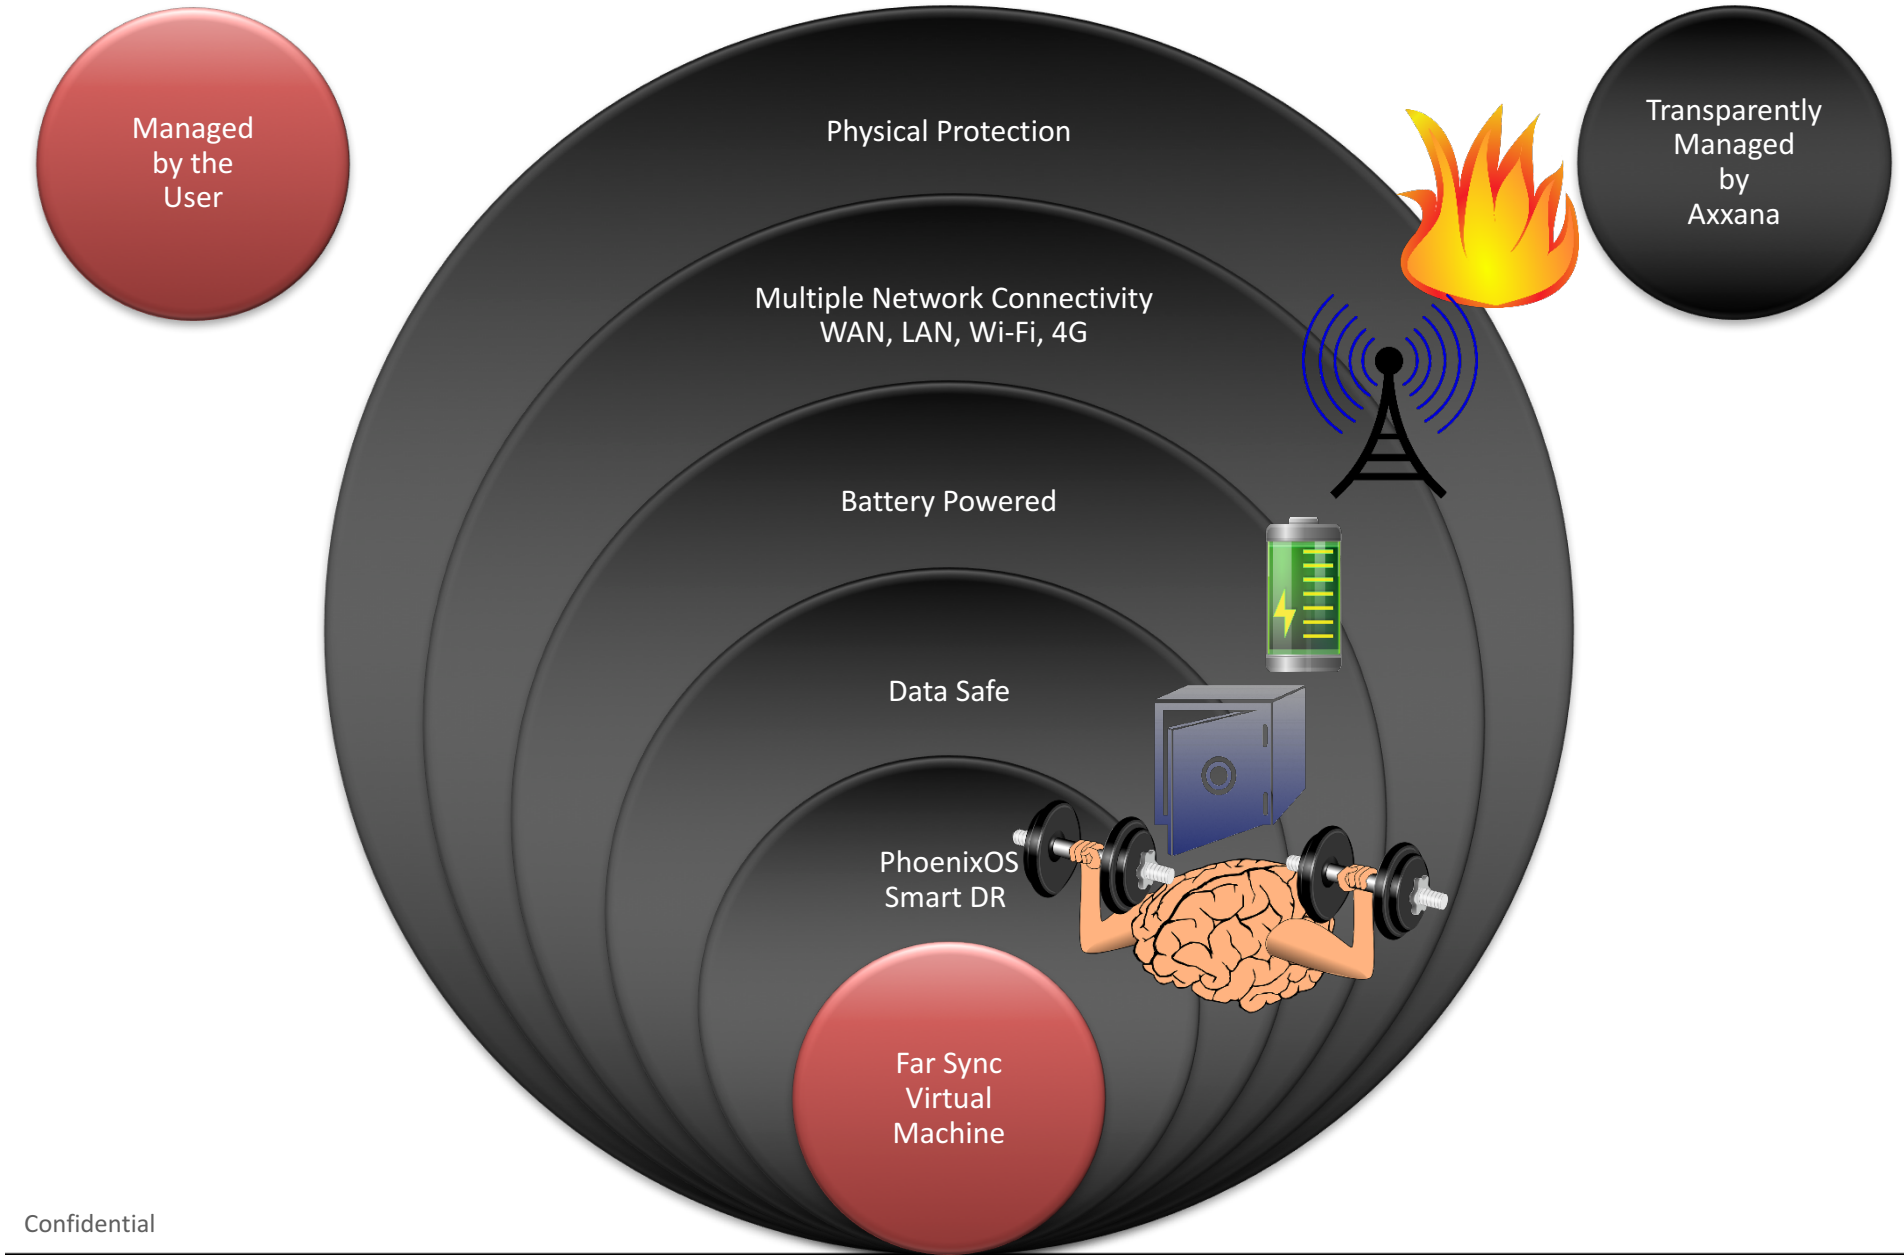

Confidential

Active Data Guard Far Sync with Axxana

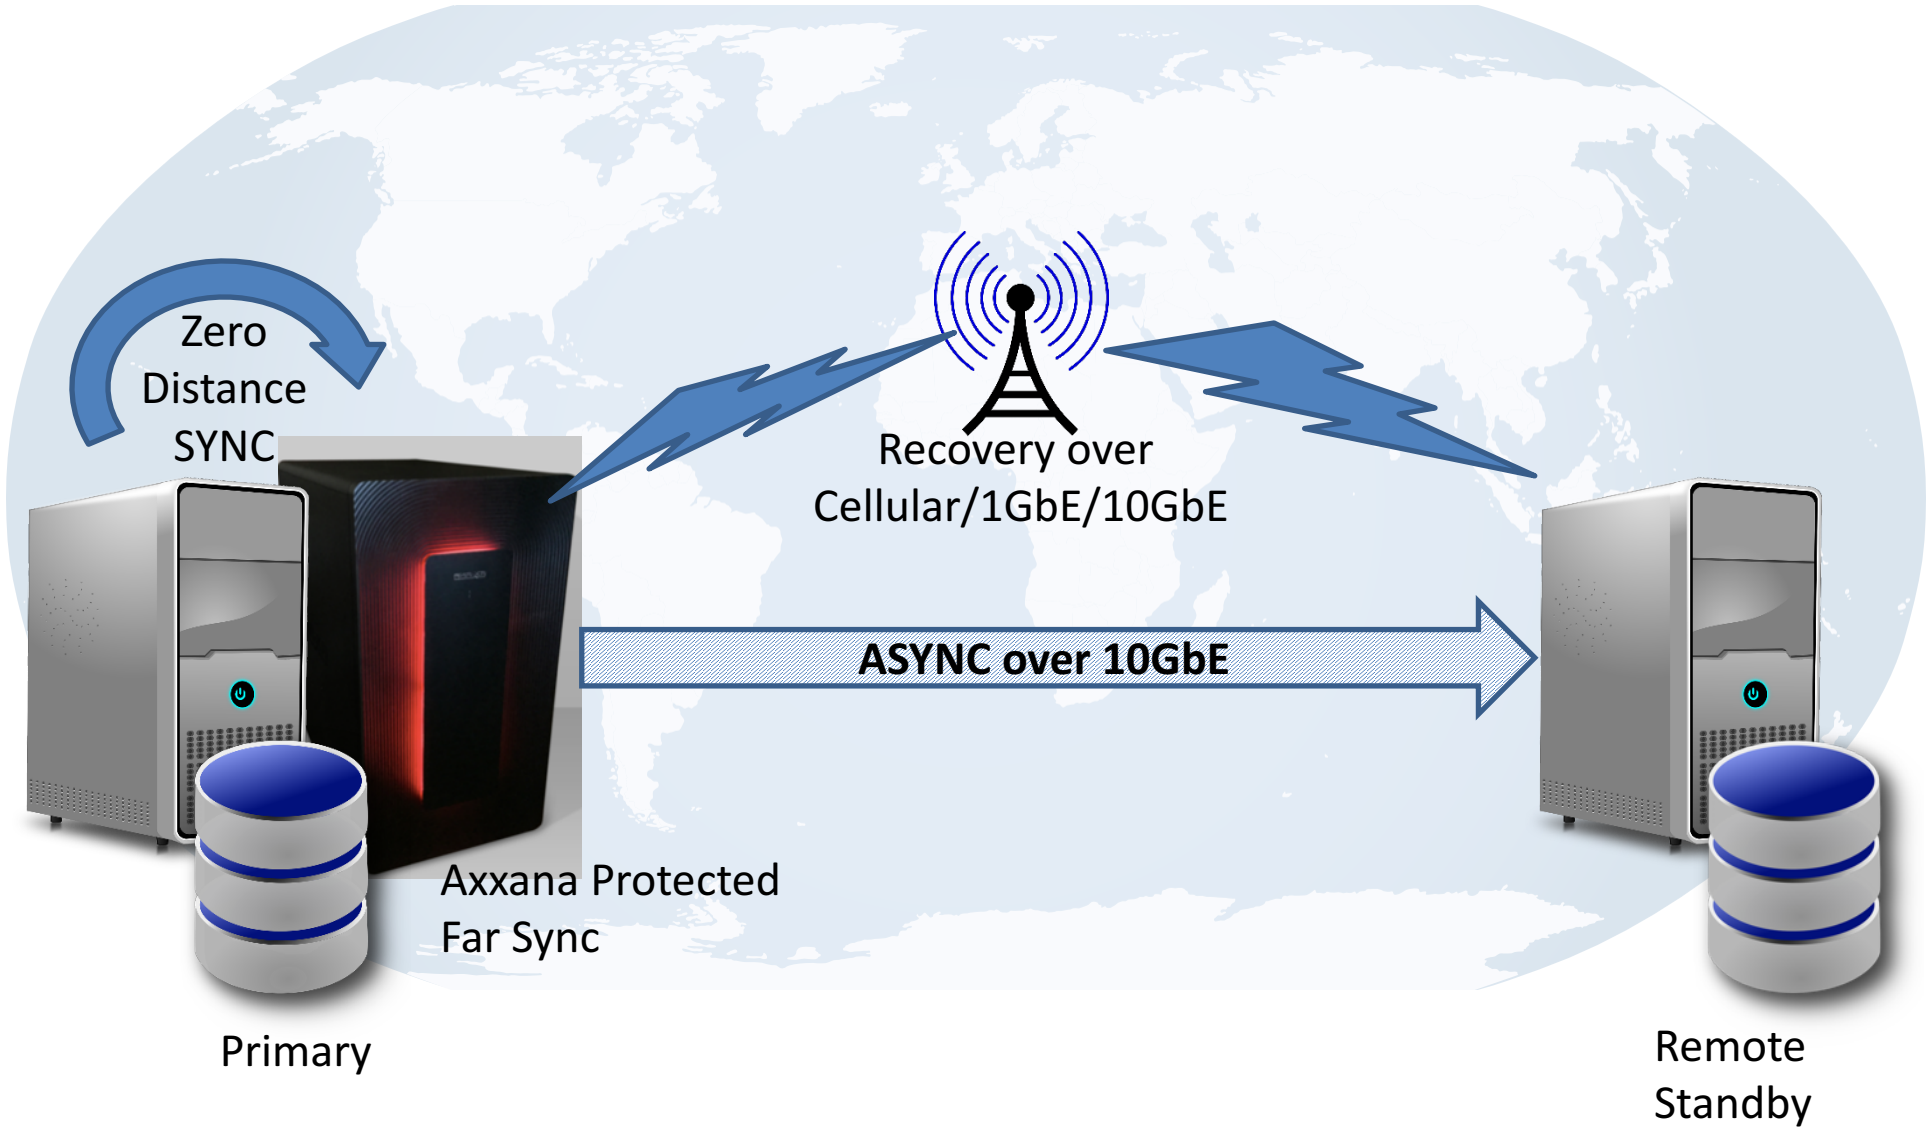

Confidential

© All Rights Reserved. Axxana 2017

What is Axxana all about ?

- A very resilient flash based storage and server in a (black) box
- Direct fire of up to 1100°C for an hour
- 250°C for 6 hours
- 5m drop on reinforced concrete
- 2.3 tons of weight
- 1 m water pressure
- Pierce force of 230 kg rod with cross-section of .25 cm^2 dropped from 3m height
- Self sufficient power source
- Independent cellular and WiFi communication

In a nutshell – a bunker in a box

Winning Combination

ORACLE®

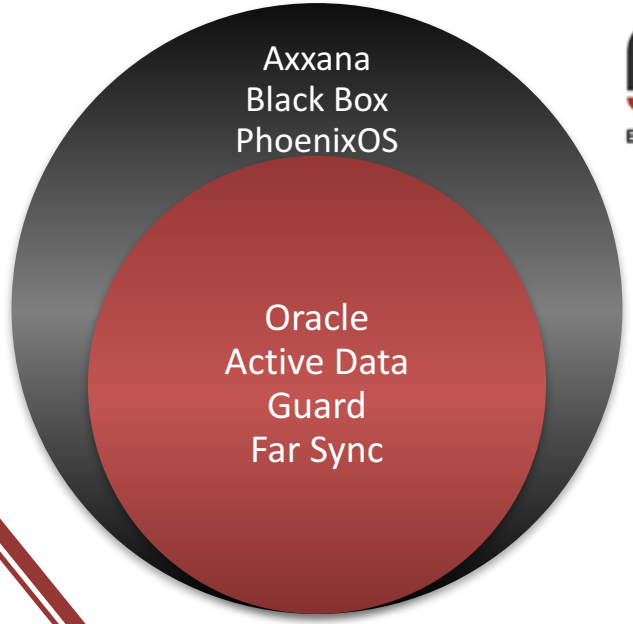

AXXANA
BUILT TO LAST

Optimized for Far Sync

- ✅ Zero Data Loss at any distance
- ✅ No RTT (latency) Compromise
- ✅ True Protection against Link Failure
- ✅ Application Consistency Guarantee

Confidential

© All Rights Reserved. Axxana 2017

AXXANA

BUILT TO LAST

Thank you!

Alain.Azagury@axxana.com

© All Rights Reserved. Axxana 2017