Disaster Recovery: Past, Present, and Future Challenges

Alex Winokur
CTO
alex@Axxana.com

AXANA

BUILT TO LAST

The Past

Setting the Seen

- 1934 – First UPS (patent grant)
- 1952 – First commercial computers
- 1952 – First tape drive
- 1956 – First disk drive
- 1964 – First mass scale commercial computer (IBM 360)
- 1970 – Clustered system for FAA (IBM 360)
- 1974 – X.25 first commercial data communication network

DR Chronology

- 70th
  - Organizations began to be totally dependent on computers
  - Terror in major industrial countries – Germany and Italy
  - With the lack of good communication, efforts were concentrated around HA
  - Critical applications were run twice
- 80th
  - First regulation requiring banks to have a testable backup plans
  - Continuous evolution of communication and clustering technologies (DEC)
- 90th
  - Backup windows issues laid the foundation for snapshots in form of IBM T0 copy
  - February 26, 1993 - World Trade Center bombing. Introduction of SRDF (separation data from applications)
  - Shaping of the DR solutions as we know today

DR Chronology

- 2000th
  - 11 September 2001
  - New SEC regulations “Interagency Paper on Sound Practices to Strengthen the Resilience of the U.S Financial Systems”.
  - Deduplication and Backup to disk
  - Standardization of DR solutions
- 2010th
  - Virtualization
  - Active active
  - Cloud and DRaaS

AXANA

BUILT TO LAST

The Present

DR from Another Angle

- The nature of disasters
- Configuration limitations
- The misconceptions and realities of DR
- Understanding data

What is Disaster Recovery?

- A disaster is an inability to continue operation at a primary site
- If data is lost, then no equipment at he primary site can be used to recover it
  - Fire
  - Flood
  - Terror
  - Earthquake
  - Weather hazards
  - Miscellaneous (storage, UPS, cooling, etc…)
- Disaster recovery – A procedure to resume operations by failing over to secondary site

Two Major Types of Disasters

- Regional/Nation-wide disasters
  - Due to events which become publicly advertised
    - Weather hazards
    - Blackouts
    - Earthquake
    - Floods
- Local disasters
  - Due to events not disclosed to the public
    - Local power failures
    - Local fires and floods
    - Miscellaneous failures like failure of cooling system, storage, communications, etc. (the majority of disasters fall in this category)

A Catastrophe!

A disastrous event in which the DR procedures dramatically miss their SLA

Some Examples

- September 11 in Manhattan, primary in twin tower, secondary in New Jersey
  - DR plan – Secondary site failover
  - Inability to move key personnel to secondary site
- 2003 blackout, sync mirror replication
  - DR plan - Generators up
  - No water supply, cooling system fails – both sites down
- Power failure in major snow storm in New England, sync mirror replication
  - DR plan – Secondary site failover
  - Power outage at the central communication switch just before power outage – Sync replication stopped, results in data loss

Typical B.C. Implementation

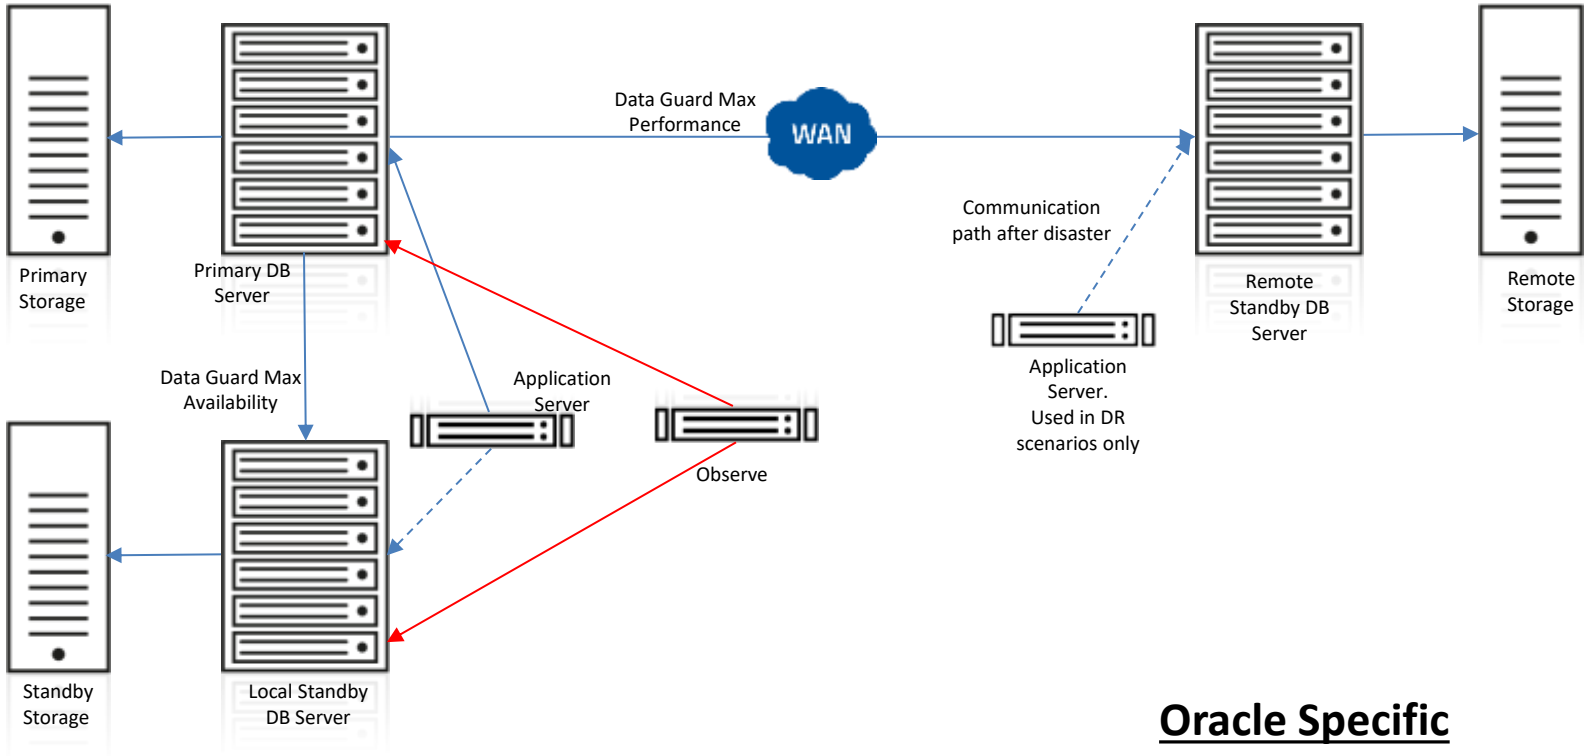

Oracle Specific

FSFO	Fast Start Fail Over	Database server failover
TAF	Transparent Application Failover	Application server failover
TG	Transaction Guard	State of in-flight transaction

The Utopia

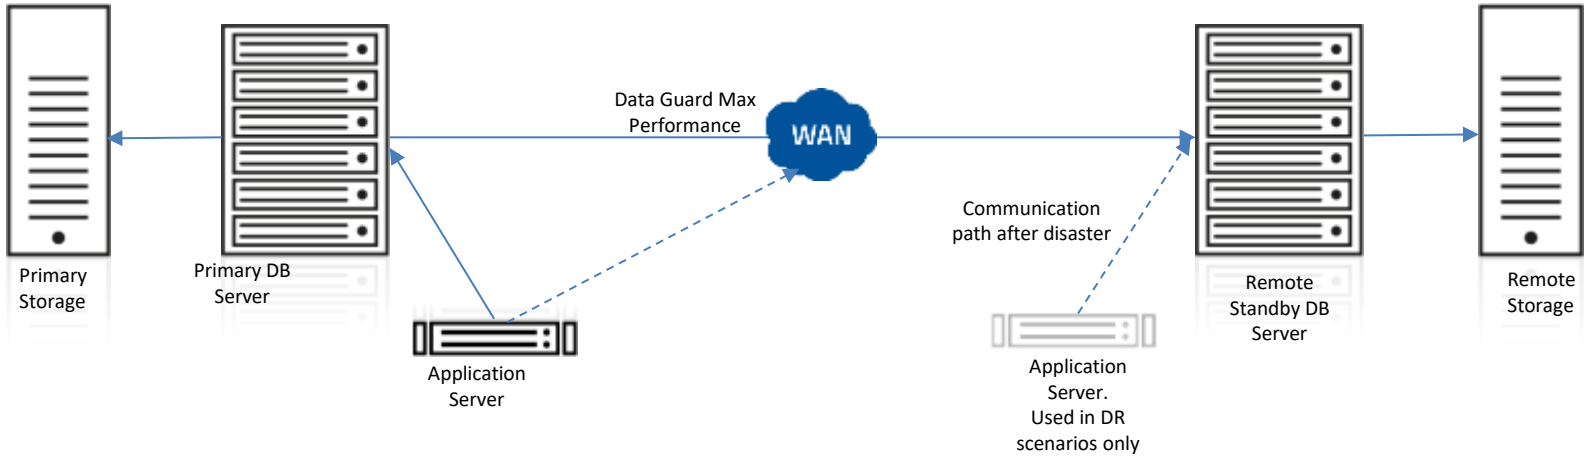

- Seamless failover over unlimited distance
- Active-active environment with load balancing capabilities
- Partially possible with a major application rewrite

A tradeoff between survivability and performance

The Five Most Common Misconceptions About DR

- I can anticipate all disaster scenarios
  - Too much dependence on external infrastructure
  - Murphy's low is king
- Sync Replication implies no data loss
  - The issue of rolling disaster
- If I succeed in DR tests, I am ready for the real thing
  - Switch over is not the same as fail over
  - No space for unexpected events
- RTO of 4 hours .means I am operational in 4 hours after a disaster
  - From which point in time? Disaster time? Start failover time?
- RPO=5 minutes means I can lose 5 minutes of data
  - It really means: one must be able to reconcile the last 5 minutes of data

Understanding Data

- Real world data
- Historical data – warehousing
- Unstructured data (Office applications)
- Temporal data
- Personal data

Let's Concentrate on the Real-World Data

- There are always at least two instances of the data
  - In the database
  - In the real world
- The exitance of these two instances is the cause of most "data loss" evil
  - Causes inconsistencies and operation disruptions
  - It is an embarrassment to the organization hence loss of prestige
- On the other hand – These two instances are in the core of all the data reconciliation techniques

Some Examples

- The application - a subscription to an online service
- A disaster while a customer is canceling a subscription
  - The cancelation request is lost
- Two scenarios
  - The customer gets a confirmation (eMail) but continues to be billed
  - The customer does not get confirmation (eMail) but continues to be billed
- When the customer complains, in case II the company can always claim that he never canceled his subscription!!!

AXANA
BUILT TO LAST

Two Costly Options When Data is Lost

- Data reconciliation – aligning the systems back with the real world
  - Tedious, error prone (mostly manual) process of data discovery and reentry
  - Involves interaction with vendors, customers and other system users
  - Long recovery time even for few lost transactions
- Resumption with data loss
  - Cost of unexpected operation disruptions when inconsistencies are finally detected
  - Cost of the lost transactions themselves
  - High legal and compliance costs
- In both cases enterprise's reputation and prestige is greatly compromised

Some Examples

- Inventory system in a car manufacturing plant
  - After a disaster real inventory does not match what is in the inventory database
  - Option one do a recount
  - Option two do nothing
    - At some point manufacturing will have to stop due to shortage of paint for example
- CRM ordering system
  - After a disaster some orders are lost
  - Option I publish an email asking every customer to resubmit last day orders
  - Option II do nothing and wait for complains

Different Value of Same Data Type

- Losing telephone billing records
  - The customer is very happy
  - If loss containable no need for reciliation
- Losing law firm billing records
  - Customer very unhappy. If my billing records were lost, what other important data was lost too?!?
  - Reliability of the firm is compromised
  - Billing data should be protected at all cost

AXANA

BUILT TO LAST

The Future

Ever since the fall of the holy temple, it has come to be that prophecy was taken away from the prophets, and given to the infants and the fools

What There is Ahead For Us?

A continues race between threats and technology

The Threats

- The facts
  - Amount of data is growing exponentially
  - Global worming
  - Cyber attacks
  - Systems are becoming more and more complicated
- The speculations
  - Systems are becoming less reliable
  - Competent IT skills are becoming scares
- The outcome
  - Disasters will become more frequent with more sever consequences

The New technologies

- Simplify Infrastructure
  - Cloud
  - Converged systems
- Automation and standardizations are key
  - Very strict standards for infrastructure layout and programing style
- Deep learning and AI
  - Will there be enough data?
- Identify the most critical data and don’t try to save $ on its protection
  - The problem is the ROI model

The Economics of DR Solutions

- Until a disaster occurs the investment is useless – just like insurance
- On the other hand – Can we find ways to utilize the spare equipment more economically?
  - Backups
  - Data mining
  - Reporting
  - Part of elastic cloud?

Food for Thought

- Can there be an insurance model?
  - How to calculate coverage?
  - How to calculate Indemnification?
  - How to calculate premium?

Time	Probability of disastrous event
1 year	5%
2 years	9.8%
3 years	14.3%
4 years	18.5%
5 years	22.6%

AXANA

BUILT TO LAST

Thank you!