Data Guard Fast-Start Failover

At Amazon.com

November 14, 2007

John Smiley

amazon.com

Goals for Fast-Start Failover

- Reduce time to failover
- Simplify failover process
- Set standards for Fast-Start Failover deployment
- Reduce costs

amazon.com

Implementation

- Observer
  - Multiple hosts
  - Wallet for SYS password
  - Wrapper to start/restart automatically
  - Easily relocated
- Monitor
  - Incorporated into existing monitoring infrastructure
  - Alerts if FSFO readiness is compromised
  - Warns if flashback database history is insufficient to reinstate
- DB_ROLE_CHANGE trigger
  - Directory service update
- Configuration verifier
- Autostop Script

amazon.com

Failover

- Only the observer should initiate failover
- Database failover is not enough – need client failover
- Re-establish connectivity
  - Service relocation or
  - Change directory service to point to new primary location
- Notify clients
- Reconnect / retry

amazon.com

Experience

- Application failover time =
  database failover time + directory service propagation time
- Reliable
  - Always have a good primary after a failover
  - No split brain conditions
  - Data integrity maintained
- Fast, automatic standby reinstatement
  - Flashback restore time depends upon the number of distinct blocks changed during the 30 minutes prior to failover
  - Much faster than RMAN or other restore methods
  - Recovery time depends upon amount of redo generated between the restore SCN and the standby_became_primary_scn

amazon.com

Experience

- Flashback database storage requirements
  - v$flashback_database_stat estimate can be inaccurate
  - Flashback database retention target
  - Number of distinct blocks changed
  - Measure peak generation rate for the required retention period
- Maximum Availability Mode
  - Increased commit latency
  - Small percentage increase in typical application transaction time
  - Decreased throughput can be regained by increasing parallelism
  - DML pause for LGWR NetTimeout duration
  - DML may see short delay when mounting/dismounting the standby

amazon.com

Experience

- No application changes required
- Use cases
  - Instance failure
  - Host failure
  - Network failure
  - Fleet migration
- Test!
  - Only reliable if the configuration is correct
  - Verify configuration and procedures by performing failovers

amazon.com

Experience

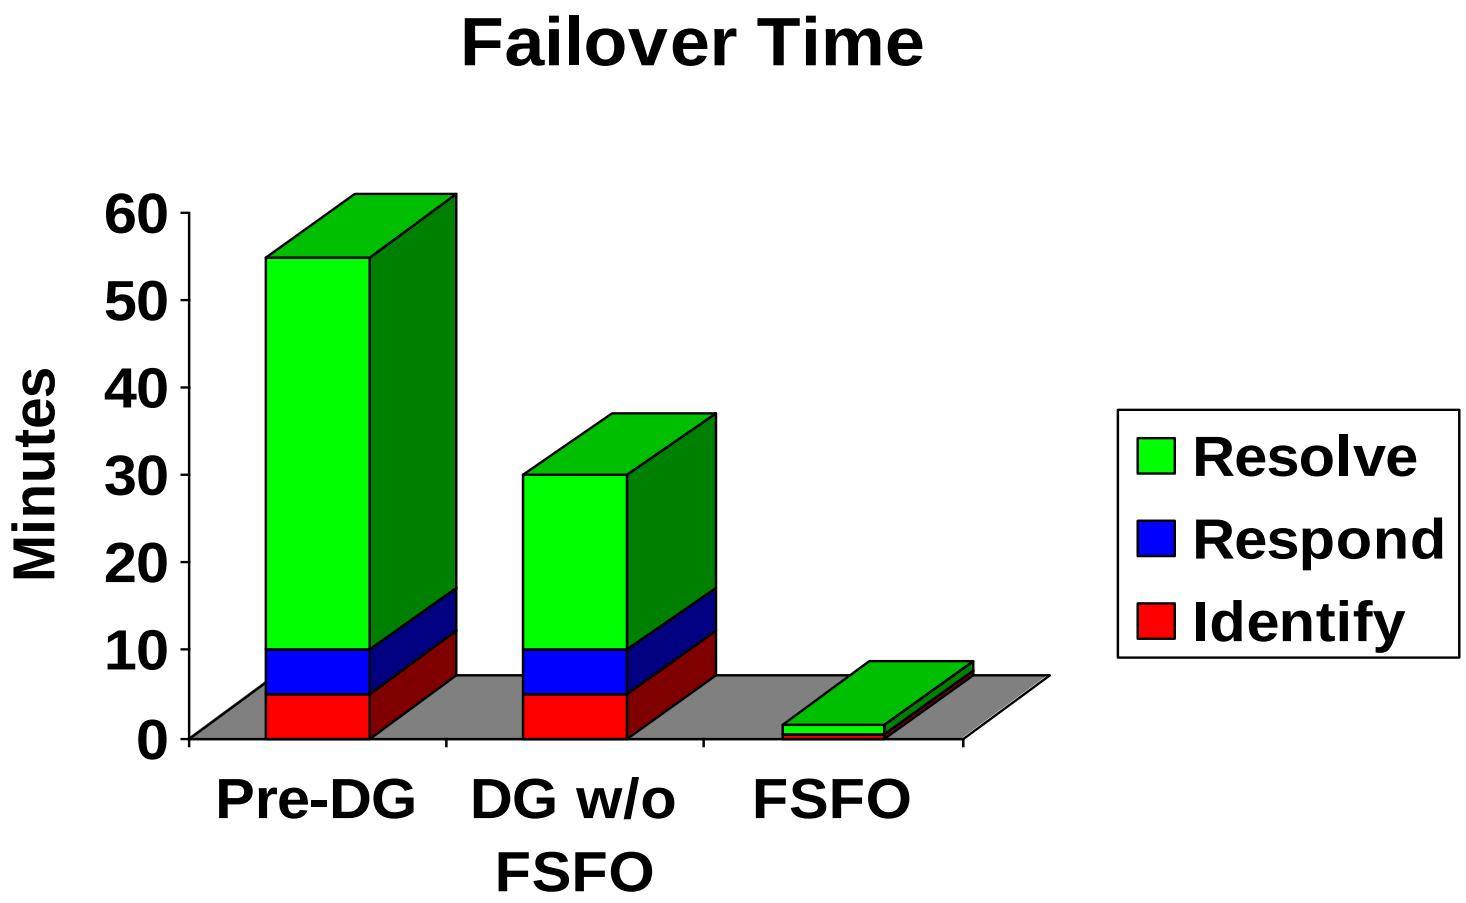

amazon.com

Other Possible Uses

- Use inexpensive commodity hardware
  - Less durable
  - Compensate with more standbys
  - Overall TCO lowered
- 11g custom failover conditions
  - Brownouts
- 11g Maximum Performance Mode FSFO
  - Asynchronous redo transfer
  - No impact to commit/transaction latency
  - Disaster recovery over long distances

amazon.com