SCHOLAR
IT SOLUTIONS

SQL SERVER AG

mailto:Sreelekha@scholaritinc.com

Domain-Independent AGs

SCHOLAR
IT SOLUTIONS

ROHIT KUMAR
DATABASE OPERATION MANAGER

LinkedIn: https://www.linkedin.com/in/rohitdba/
Twitter: @RohitKu2311

Over 11 years of experience

- SQL Server 2000 to now
- PostgreSQL
- MySQL
- MongoDB
- Cassandra
- Azure SQL
- AWS RDS

Certifications:

- Microsoft Certified Trainer
- MCSA
- Azure Solution Architect
- Azure Administrator
- Azure Database Administrator
- AWS CCP
- Cassandra Administrator

www.scholaritinc.com
Pleasanton,CA,USA
For More Information Please Contact
+1 (925) 999-0678
Sreelekha@scholaritinc.com

RPO & RTO

- RPO
  - How much data loss is acceptable in the event of a failure.
- RTO
  - The maximum amount of time an application can be down before recovery is complete and users can reconnect.

ALWAYS ON???

- Always On is a collection of high availability and disaster recovery functionality with the goal to minimize recovery point objective (RPO) and recovery time objective (RTO) further below the times that can be already achieved using other existing HA/DR solutions.
- Automatic and manual failover
- Multiple secondaries
- Active secondaries
- Fast failover (Less than 30 seconds but depends on work load and configuration)
- Automatic Page Repair

ALWAYS ON TERMINOLOGY

- Primary
- Secondary
- Availability Group
- Availability Replica
- Availability Database
- Availability Group Listener
- Availability Mode
- Failover Mode
- Log Cache
- SQL Server FCI

ALWAYS ON EVOLUTION

- 2012
  - 5 replica (1+4)
  - 2 automatic failover partners
  - 3 sync partners
- 2014
  - 8 replicas
  - Azure replica
  - 2 automatic failover partners
  - 3 sync partners (1+2)
  - Introduction of Read intent feature
  - Improved dashboard with add/remove column

ALWAYS ON EVOLUTION

- 2016
  - Improved in log transport performance
  - Round Robin Load balancing in readable secondaries
  - MSDTC support
  - Domain Independent Availability Groups
  - Automatic Seeding
  - DAG – Up-to 17 readable secondaries
  - Updatable column store index support on secondary replica
  - Support of TDE
  - Basic AG support in Standard edition
    - No HA, Non-readable secondary, 1 DB per AG
    - No backup over secondary, single database failover, single replica
  - Database level health detection

ALWAYS ON EVOLUTION

- 2017
  - Read scale availability groups (cluster-less, no HA)
  - Cross platform Always on support (Linux and Docker)
- 2019
  - 6 sync replicas (1+5)
  - Kubernetes support for AG as an Orchestration layer
  - Automatic Read/Write routing to Primary

SCHOLAR
IT SOLUTIONS

DATA SYNCHRONIZATION IN ALWAYS ON

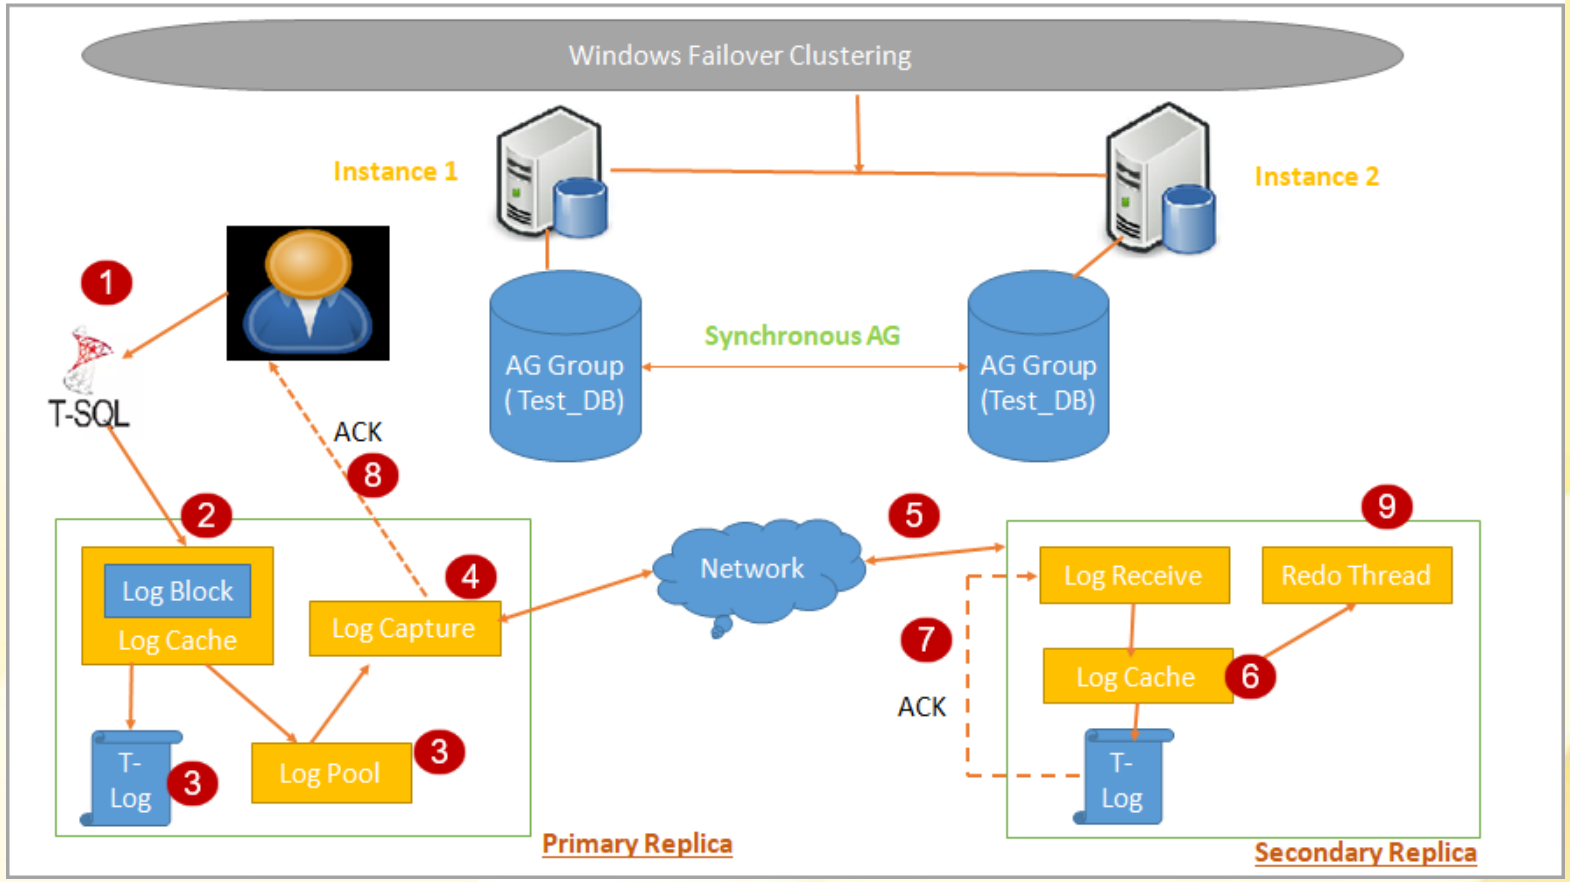

SCHOLAR
IT SOLUTIONS

DATA SYNCHRONIZATION IN ALWAYS ON

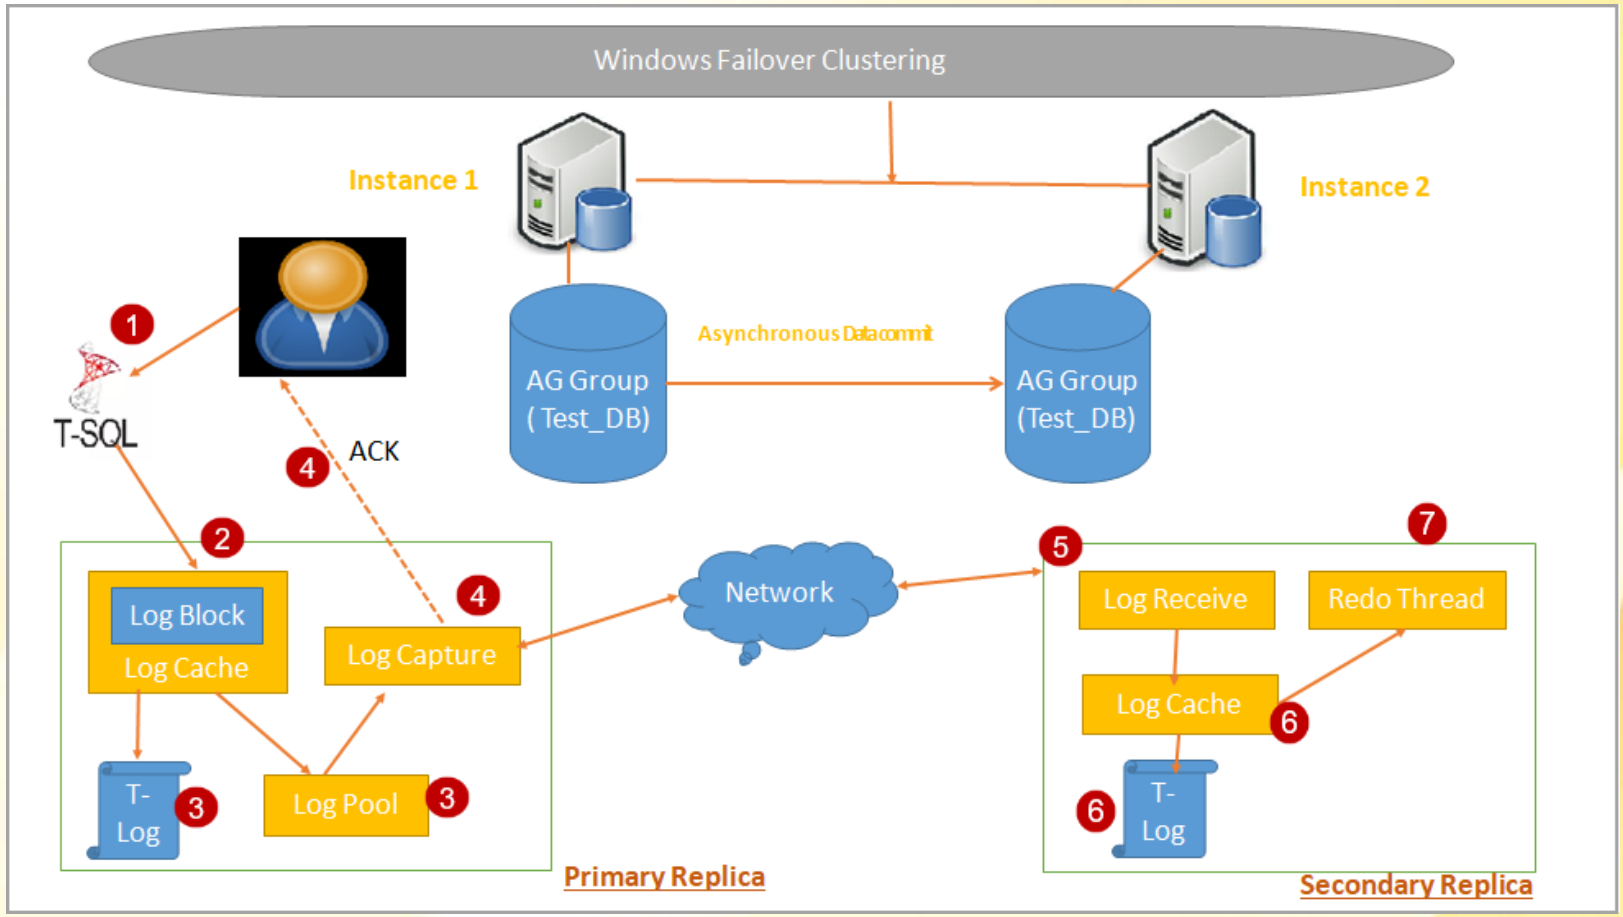

DATA SYNCHRONIZATION IN ALWAYS ON

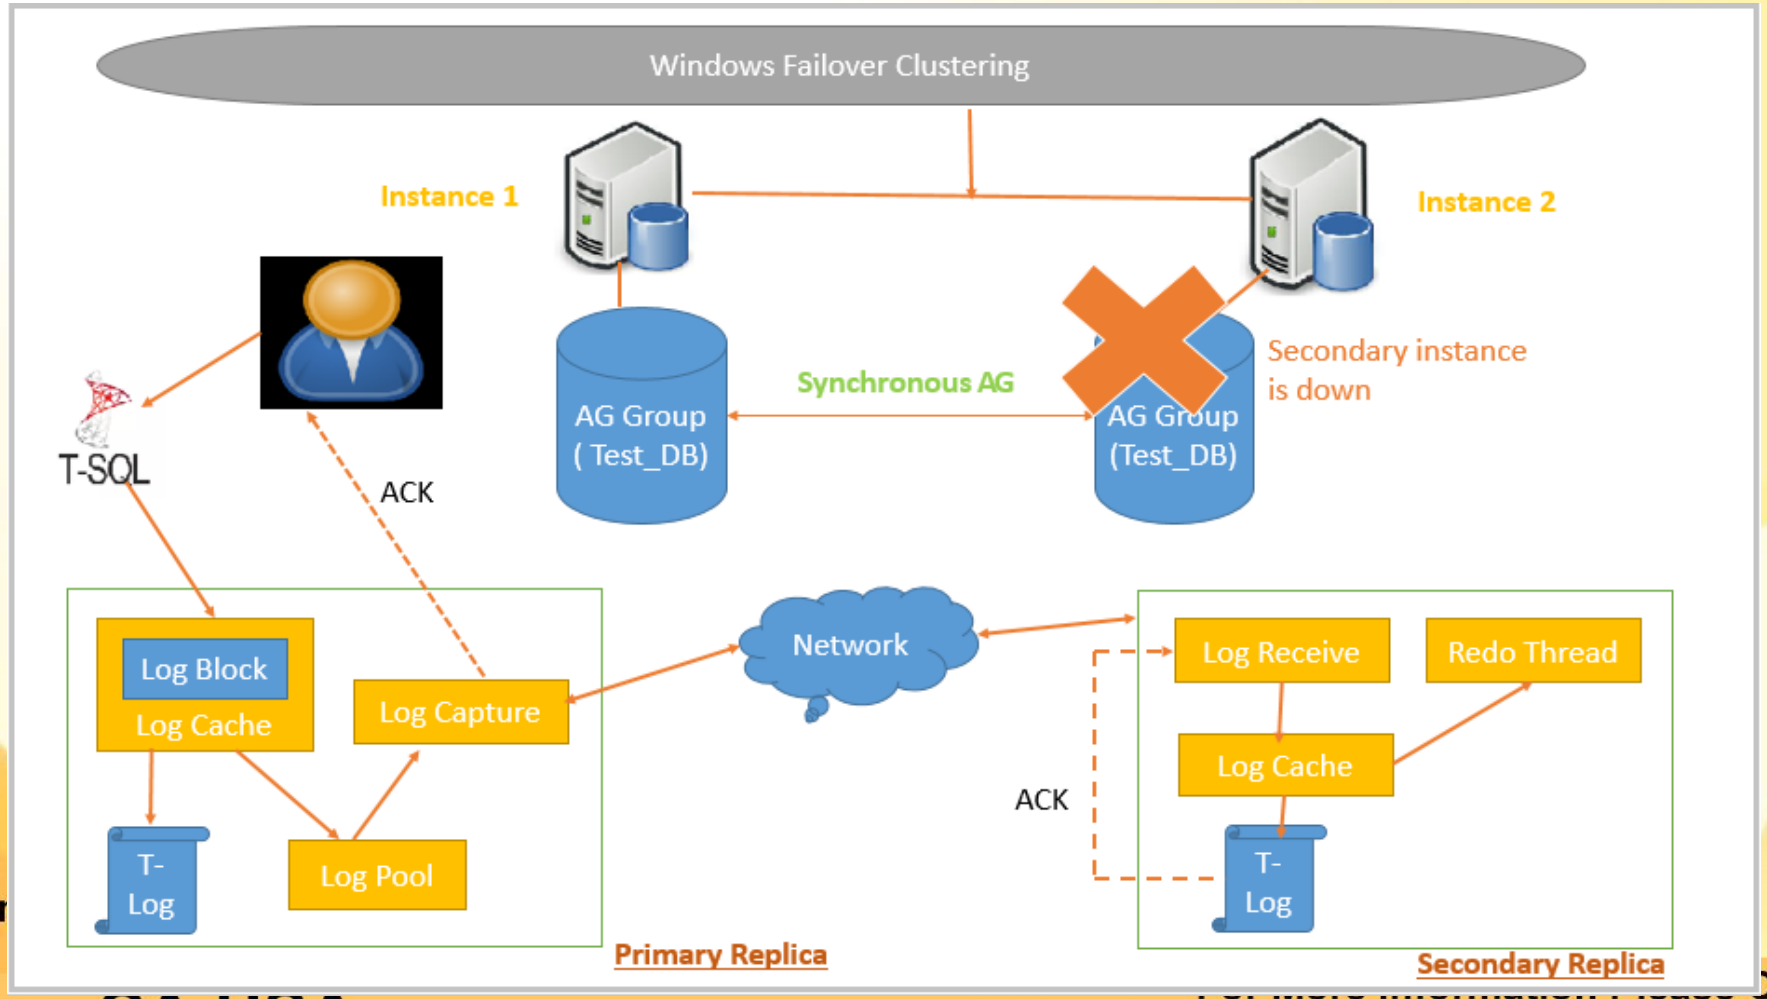

SCHOLAR
IT SOLUTIONS

DATA SYNCHRONIZATION IN ALWAYS ON

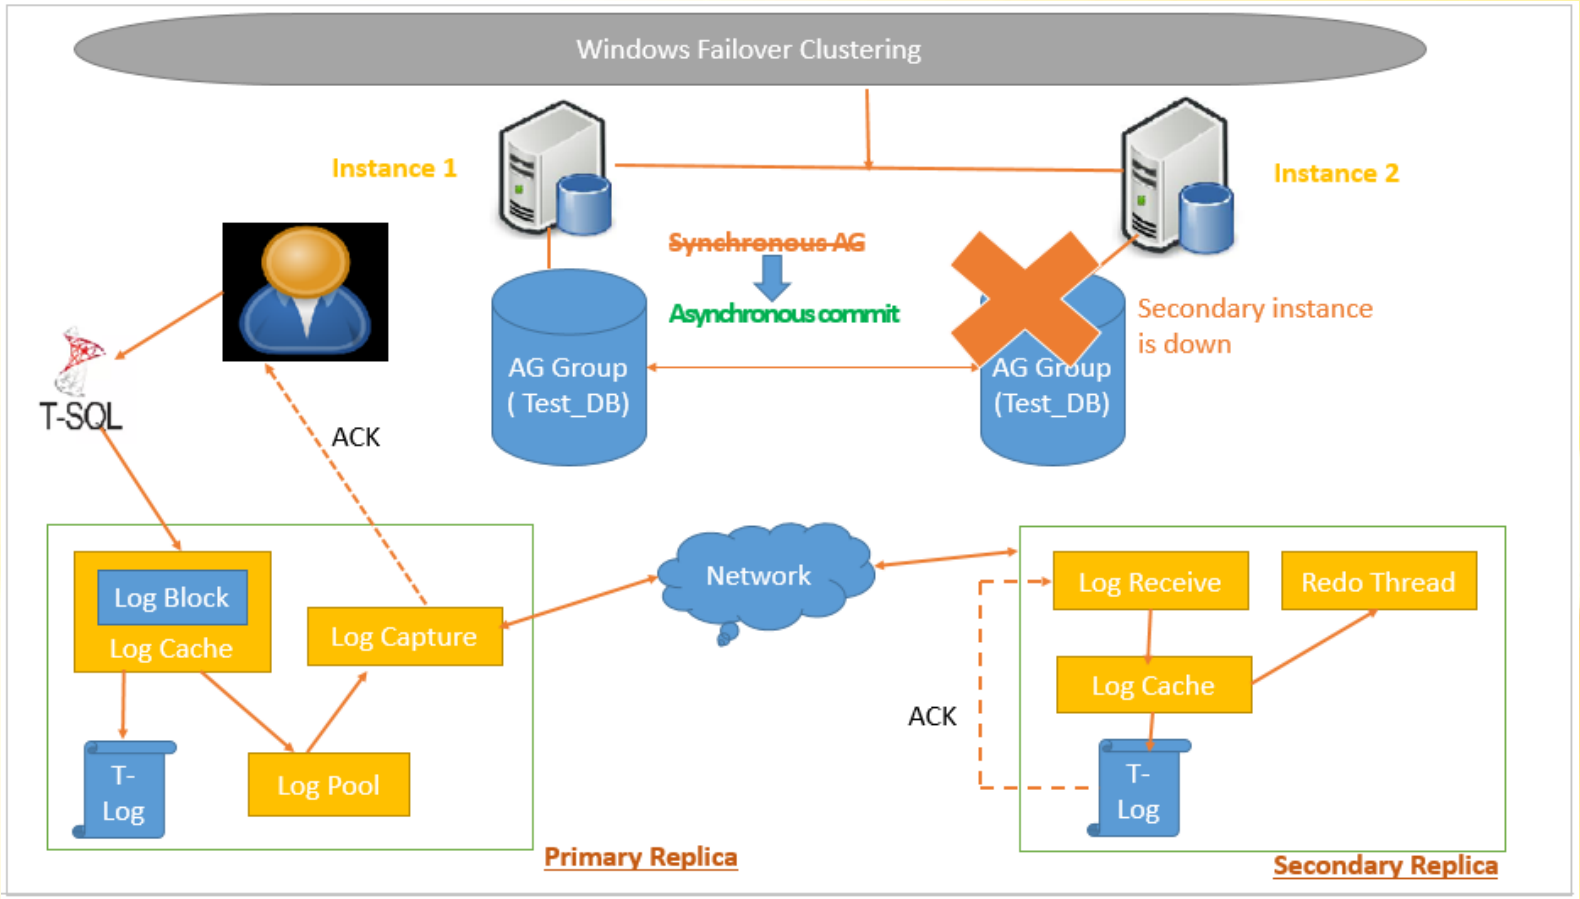

WHEN YOU CREATE AN AG

- Endpoints are created, started and properly permissioned
- Always On Health monitoring session (xEvent) is created, started.
- Availability Group is created on Primary
- Listener is created
- Secondary replica joins AG

AUTOMATIC PAGE REPAIR

- If a page becomes corrupt in a database configured as a replica in an AlwaysOn Availability Group topology, then SQL Server attempts to fix the corruption by obtaining a copy of the pages from one of the secondary replicas. Automatic page repair does not work for the following page types:
  - File Header page
  - Database Boot page
  - Allocation pages
    - GAM (Global Allocation Map)
    - SGAM (Shared Global Allocation Map)
    - PFS (Page Free Space)
- Logs the page in the MSDB.dbo.suspect_pages table

SCHOLAR IT SOLUTIONS

AUTOMATIC PAGE REPAIR (PRIMARY)

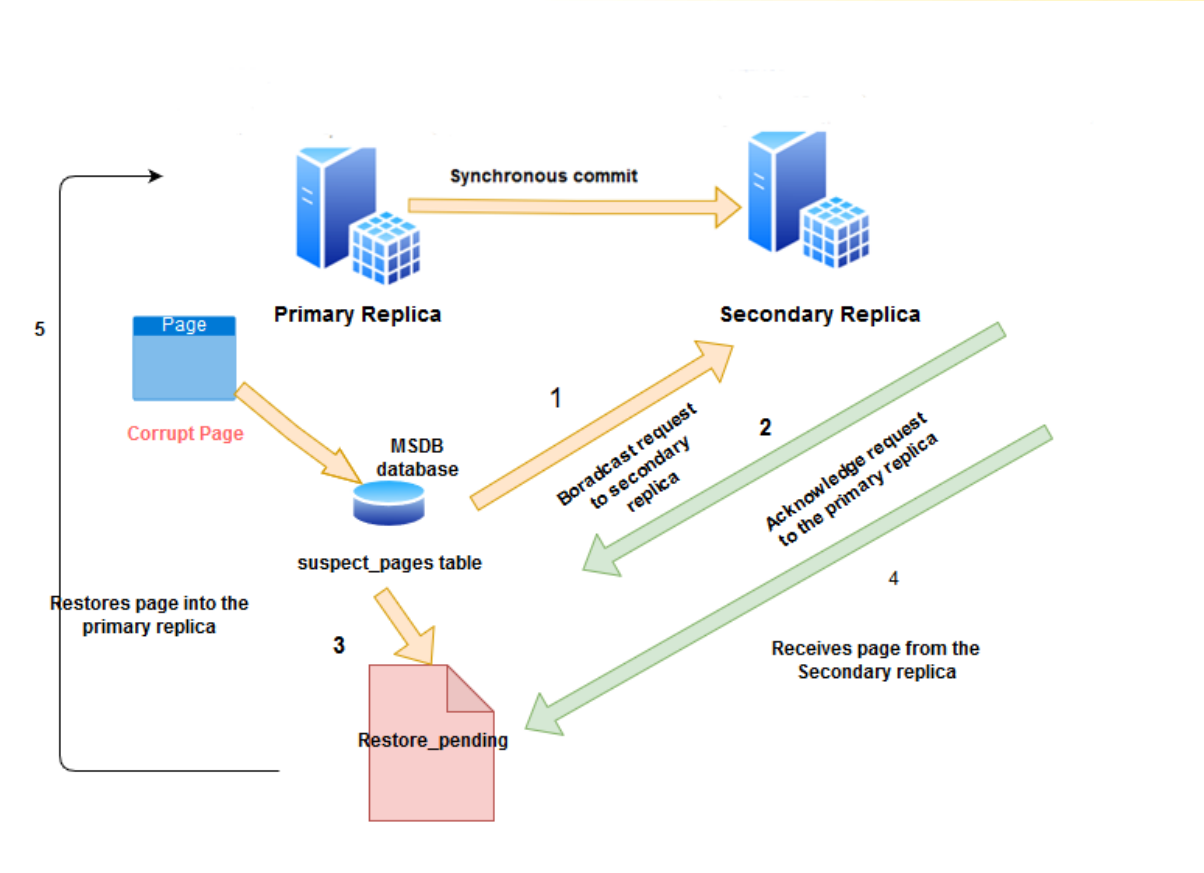

AUTOMATIC PAGE REPAIR (SECONDARY)

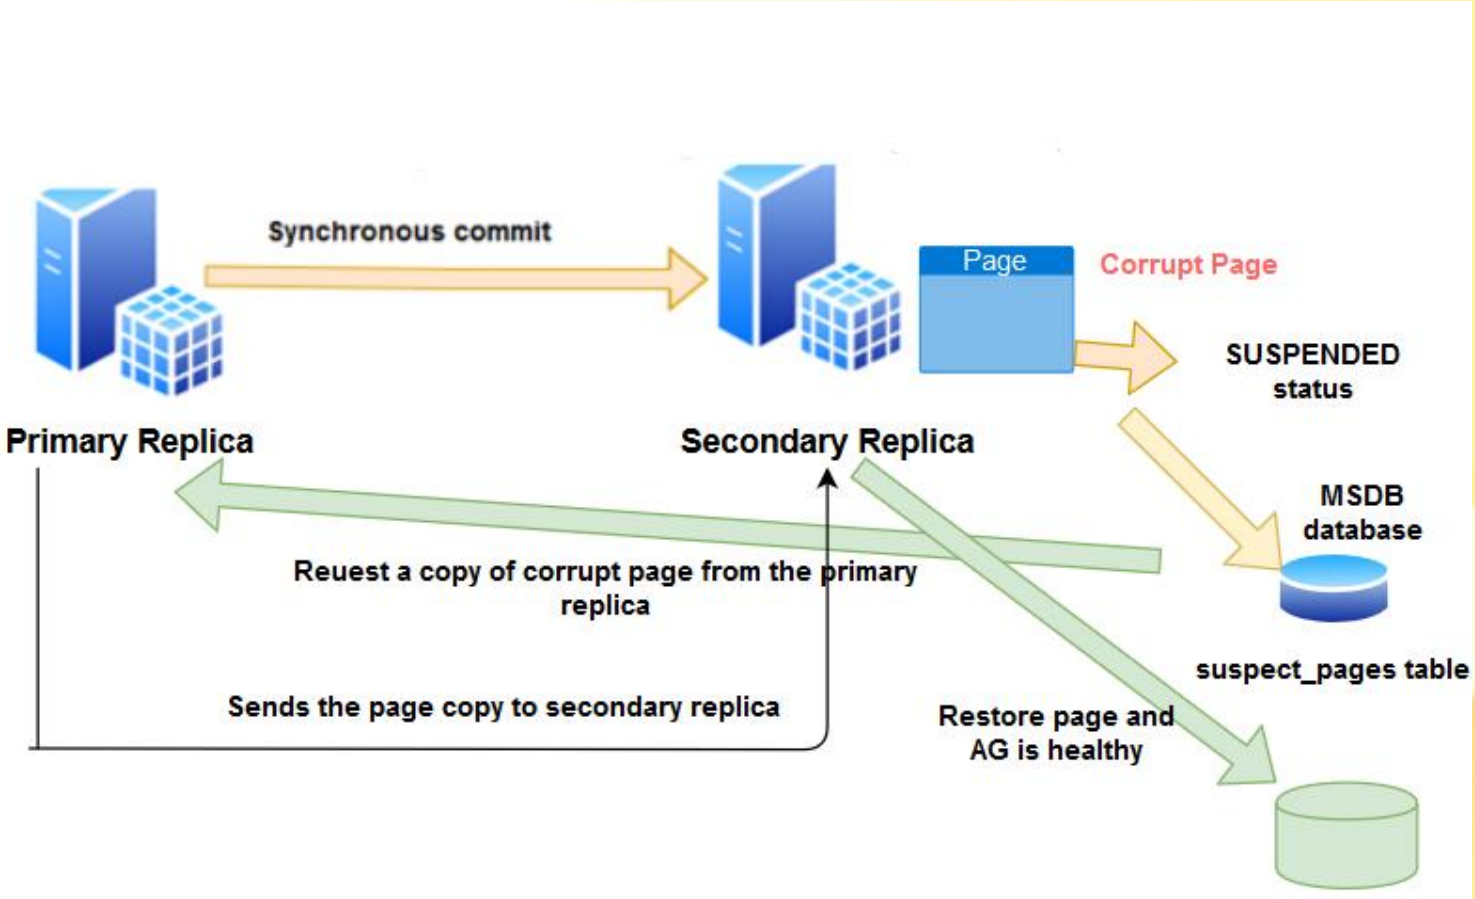

POST-CONFIG STEPS

- Failover count within a specific duration
- AG Session Timeout setting for unpredictable network
- Fail overConditionLevel/LeaseTimeout/HealthCheckTimeout

DOMAIN INDEPENDENT FAILOVER CLUSTER

- Prerequisites
  - All of the servers must be running Windows Server 2016
  - Accounts that will be used to create WSFC needs to be local admins
  - Account must have same name and passwords
  - DNS suffix for all participating nodes must be same
  - Change Remote User Account Control (UAC) LocalAccountTokenFilterPolicy registry setting.

SCHOLAR
IT SOLUTIONS

DOMAIN INDEPENDENT FAILOVER CLUSTER

- Create Login and add DNS suffix on all nodes using PS scripts
- Update host file with IP Address and DNS names on all nodes
- Disable NAT rule or adaptor with NAT rule in VM's Network section(optional/obsolet in prod environment)
- Install SQL Server and enable AlwaysOn feature
- Install Failover Cluster Feature in all nodes
- Create Failover Cluster
- Create and Backup DB
- Create Endpoint and Certificates on all Nodes
- Install Endpoint certificates on all nodes
- Create AG

SCHOLAR
IT SOLUTIONS

ALWAYS ON INTERNALS

(IMPACT OF READ WORKLOAD & ROLE OF ROW VERSIONING)
(QUERY PERFORMANCE ON READ-ONLY REPLICA)

- Read-Replica - Isolation Level
- Read Replica - Isolation Level – TEMPDB
- Read Replica - Isolation Level - Statistics - TEMPDB
- Read Replica - Fragmentation

ALWAYS ON INTERNALS

(GHOST RECORD CLEANUP)

- SQL Server has a background thread that periodically checks B-trees for ghosted records and asynchronously removes them from the leaf level of the index.
- The cleanup of ghost records on the primary replica can be blocked by transactions on the secondary replicas.
-