SQL Server Availability

SQL 2016 new innovations

Ivan Kosyakov

Technical Architect, Ph.D., http://biz-excellence.com

Microsoft Technology Center, New York

Mission-critical performance

Performance	Security	Availability	Scalability
In-Memory OLTP enhancements Greater T-SQL surface area, terabytes of memory supported, and higher number of parallel CPUs	**Always Encrypted** Sensitive data remains encrypted at all times, with ability to query	**Basic Availability Groups** With SQL 2016 Standard Edition	**Windows Server support** Support for Windows Server Core and Windows Server ReFS
Operational analytics Insights on operational data; works with In-Memory OLTP and disk-based OLTP	**Dynamic Data Masking** Real-time obfuscation of data to prevent unauthorized access	**Enhanced Always On** Distributed Availability Groups, automatic replica seeding, distributed transactions, automatic failover, load balancing, manageability	**Live migration** Faster live migration, live migration for non-clustered VMs
Query Data Store Monitored, optimized query plans	**Row-Level Security** Fine-grained access control for table rows	**Backup enhancements** Managed backup to Azure, Database Recovery Advisor	**Scalability enhancements** Hardware acceleration for TDE, parallelized decryption, TempDB optimization, and more
Temporal database support Query data as points in time	**Other enhancements** Audit success/failure of database operations TDE support for storage of In-Memory OLTP tables Enhanced auditing for OLTP with ability to track history of record changes		

Mission-critical availability

Reliable

- Detects failures reliably
- Handles multiple failures at once

Integrated

- Provides unified, simplified solution
- Streamlines deployment, management, and monitoring

Flexible

- Reuses existing investments
- Offers SAN/DAS environments

Efficient

- Allows use of HA hardware resources
- Supports Fast, transparent failover

AlwaysOn

Failover Cluster Instances
for servers

- Failover on SQL Server instance level
- Shared storage (SAN/SMB)
- Failover can take minutes based on load
- Multi-node clustering
- Passive secondary nodes

Availability Groups
for groups of databases

- Failover on database level
- Direct attached storage
- Failover takes seconds
- Multiple secondaries
- Active secondaries

Failover Cluster Instances

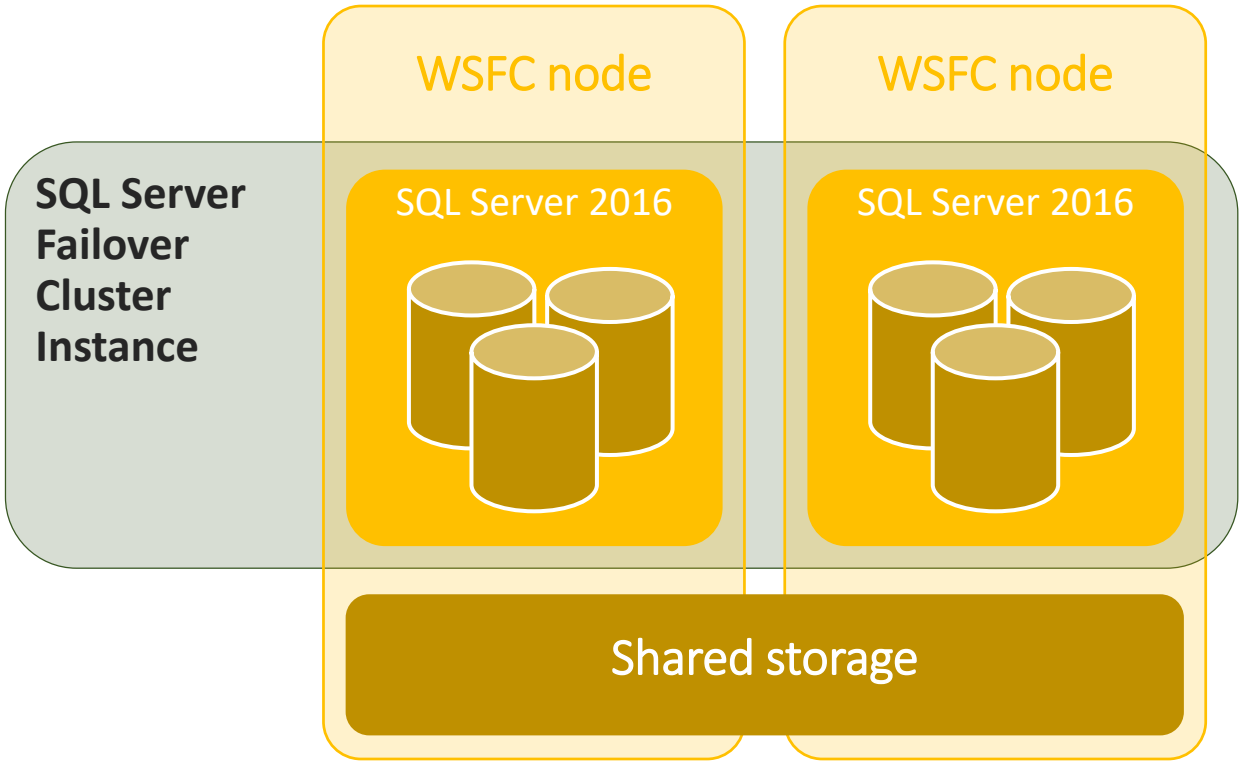

Server failover

Shared storage

Multi-node clustering

Passive secondary nodes

Failover in minutes

Availability Groups

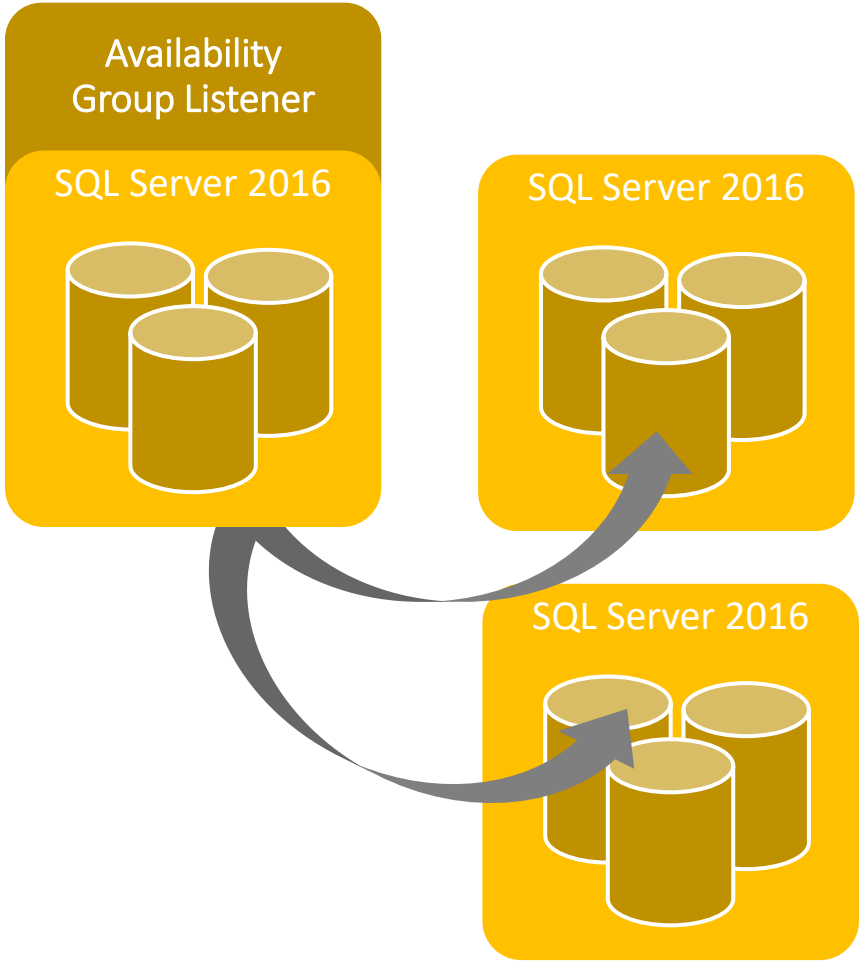

Multi-database failover

Direct attached storage

Multiple secondaries

Active secondaries

Failover in seconds

Availability Groups + Failover Clustering

Windows Server Failover Clustering (WSFC) Cluster

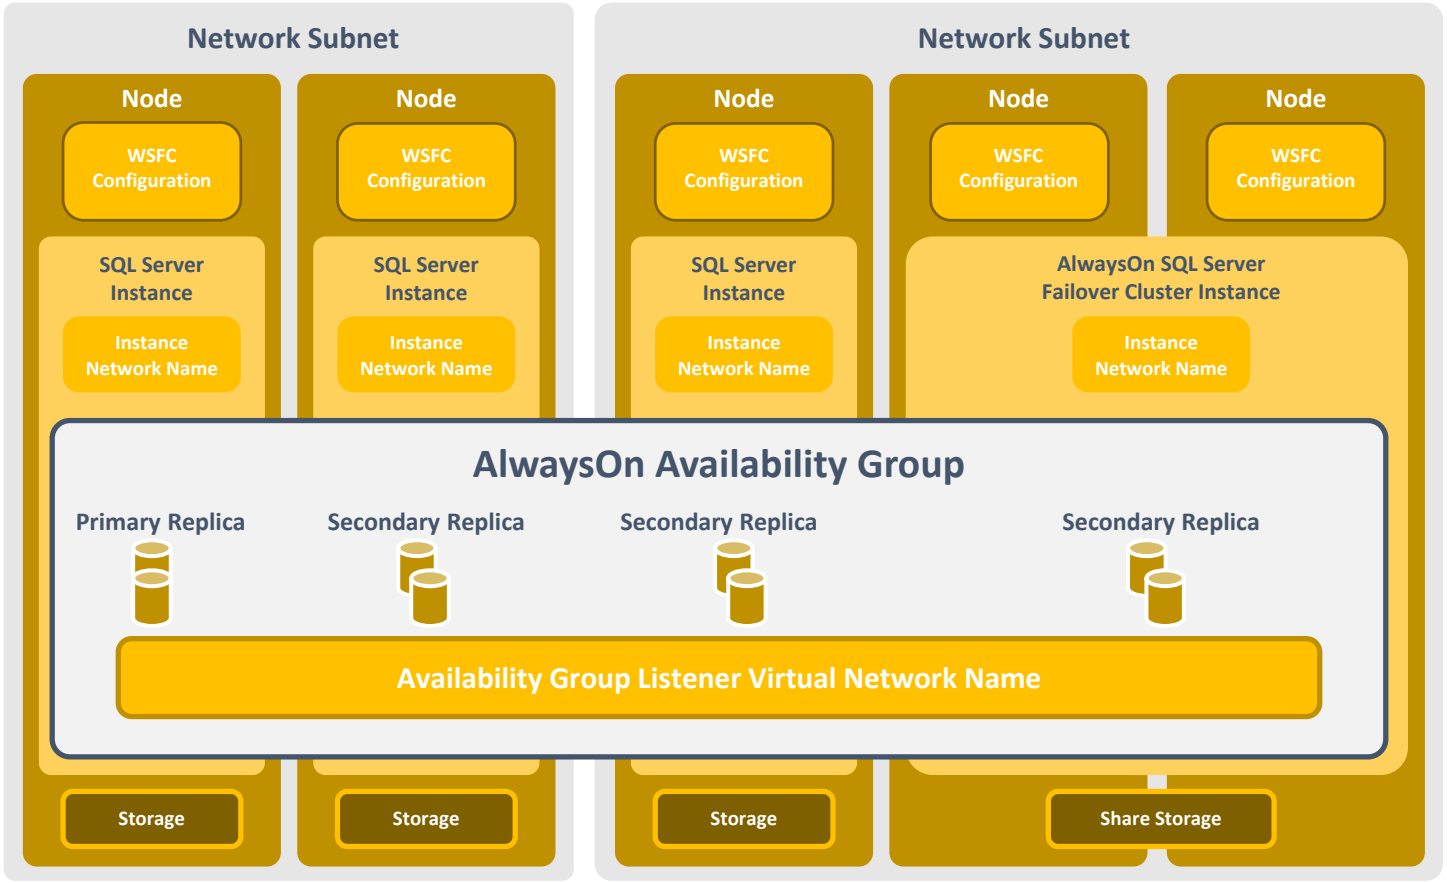

AlwaysOn:

Failover Cluster Instances and Availability Groups work together to ensure data is accessible despite failures

Basic Availability Groups

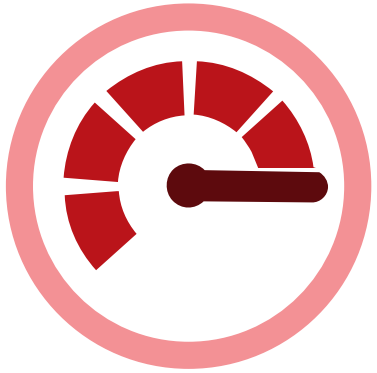

Basic Availability Groups

Available in SQL Server 2016 Standard Edition or higher

Provides failover support for single database

Replaces database mirroring feature (now deprecated)

Single replica for primary database, using either synchronous or asynchronous commit mode

Support for hybrid environments, spanning on-premises or Azure
https://msdn.microsoft.com/en-us/library/mt614935.aspx

Enhancements in Always On Availability Groups

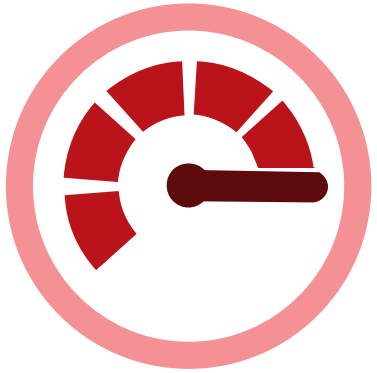

Improvements in Always On Availability Groups

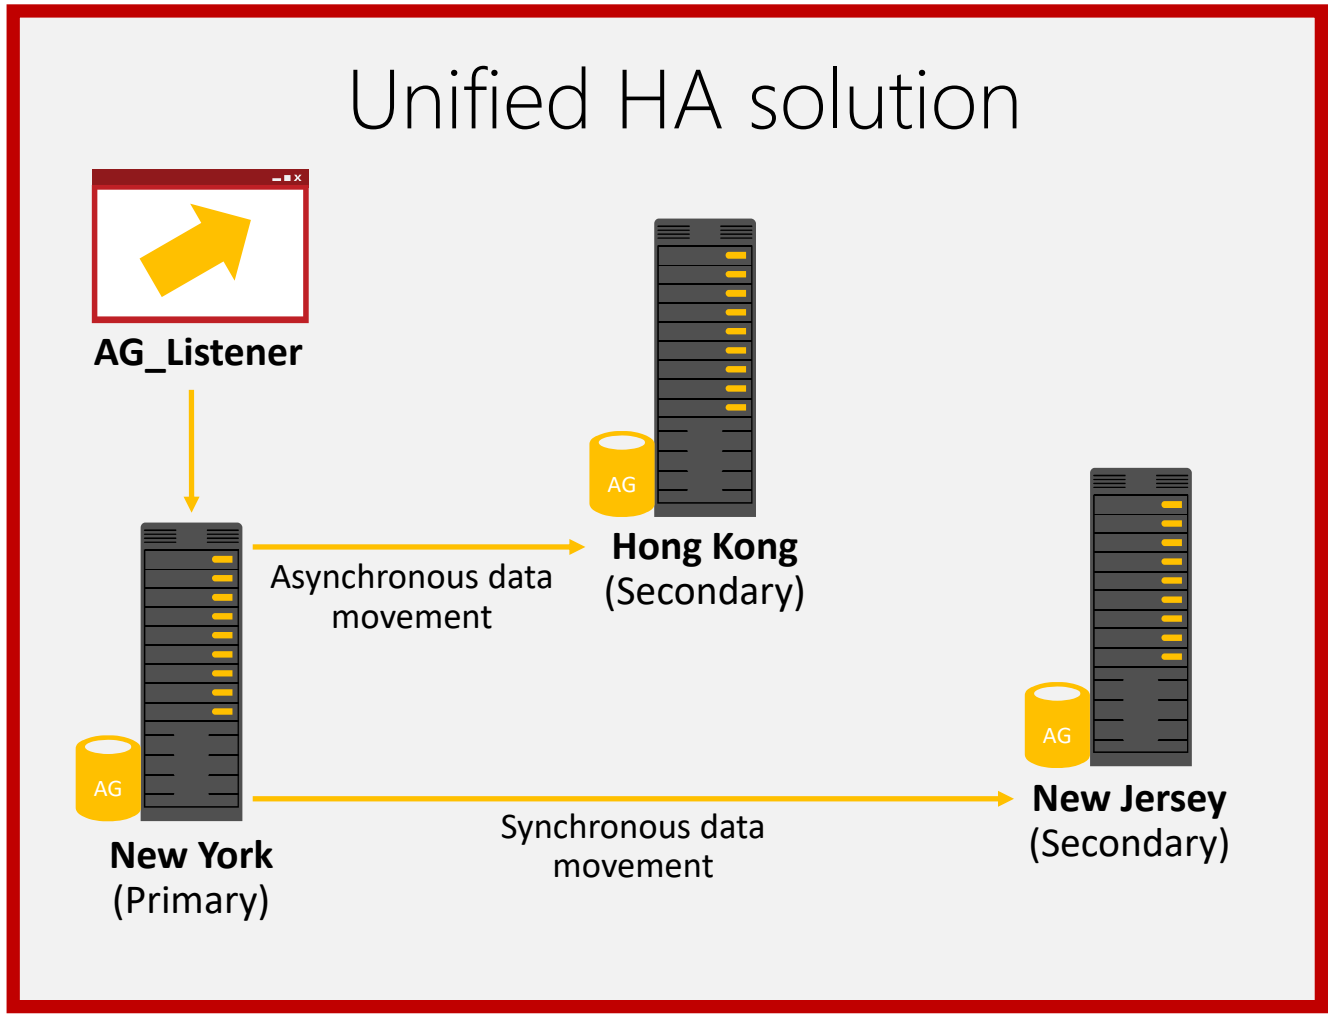

Greater scalability

Load-balancing readable secondaries

Increased number of automatic failover targets

Log transport performance

Improved manageability

DTC support with limitations

Database-level health monitoring

Group Managed Service Account

Domain-independent Availability Groups

Load balancing in readable secondaries

In SQL 2014, read-only transactions routed by the Listener went to first available secondary

Read-only routing lists can be configured to round-robin among specific set of secondaries (for each primary)

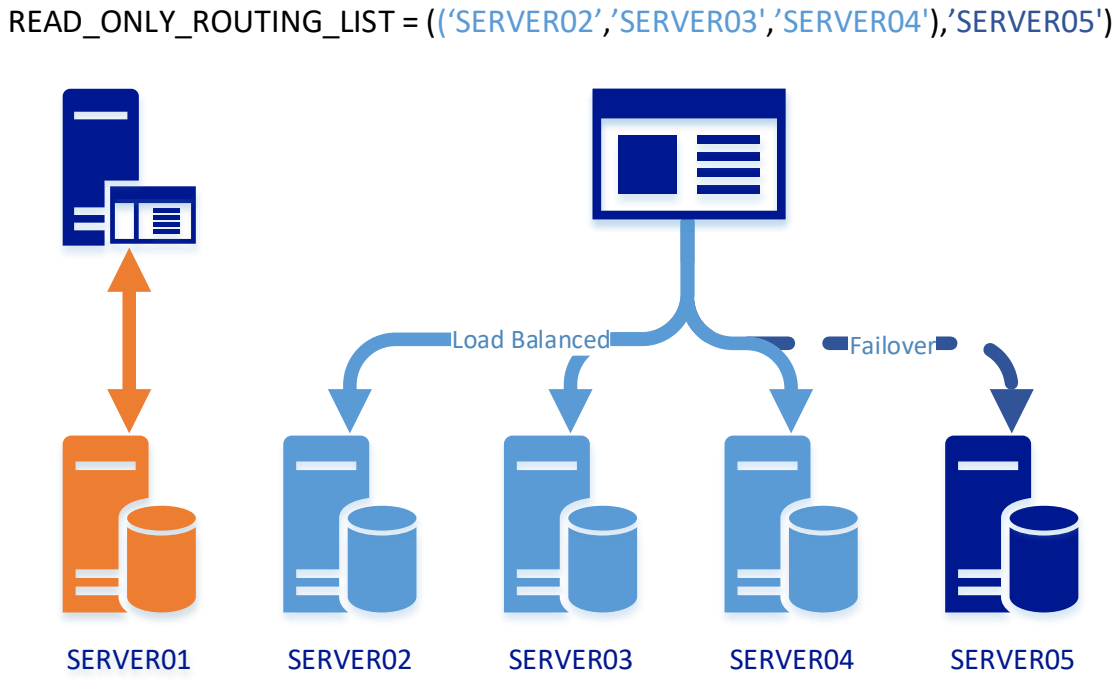

Database-level failover trigger

In SQL Server 2014, Availability Groups only monitor health of the instance

Database can be offline or corrupt, but will not trigger failover as long as instance itself is healthy

SQL Server 2016: Option to also monitor health of databases in Availability Group

Databases going offline trigger change in health status

You can now configure AlwaysOn Availability Groups for failover when database goes offline

This change requires setting **DB_FAILOVER** option to **ON** in CREATE AVAILABILITY GROUP (Transact-SQL) or ALTER AVAILABILITY GROUP (Transact-SQL) statements

gMSA support

Group Managed Service Accounts (gMSA)

- Automatically set domain scope for Managed Service Accounts
- Automatic password rotation
- Much more secure than regular domain accounts
- Enables cross-system security context

Why would I want a gMSA?

- No need to manually change passwords on all AlwaysOn instances

How does it work?

- Passwords are managed by domain

What versions will it be supported in?

- Supported in SQL Server 2014 and SQL Server 2016

Cross-Database Transactions and Distributed Transactions

Support for cross-database transactions within the same SQL Server Instance

Cross-database transactions within the same SQL Server instance are not supported for Always On Availability Groups

Support for distributed transactions

Distributed transactions are supported with Always On Availability Groups between databases hosted by two different SQL Server instances. It also applies to distributed transactions between SQL Server and another DTC-compliant server

Availability Groups must be running on Windows Server 2016 or Windows Server 2012 R2. For Windows Server 2012 R2 you must install the update in KB3090973

Availability Groups must be greated with the CREATE AVAILABILITY GROUP command and the WITH DTC_SUPPORT = PER_DB clause

Distributed transactions are not supported for database mirroring

https://msdn.microsoft.com/en-us/library/ms366279.aspx

More than two auto-failover targets

Increasing scale of solution

Increasing resiliency

Now any sync secondary can be target for automatic failover

Total of three (up from two) auto-failover targets

Domain-independent Availability Groups

Environments supported:

- Cross domains (with trust)
- Cross domains (no trust)
- No domain at all

Cluster management via PowerShell only

SQL management as normal

Use of certificate-secured endpoints like DBM

Summary: Enhanced Always On

Capability

For scalability, SQL Server 2016 adds in load balancing of readable secondaries

Increases number of auto-failover targets from two to three

Benefits

Log transport performance has been improved

Support for Distributed Transaction Coordinator (DTC): Enrolled transactions for Availability Group databases with limitations

Database-level health monitoring

gMSA: Domain-level accounts that are automatically managed

Backup enhancements

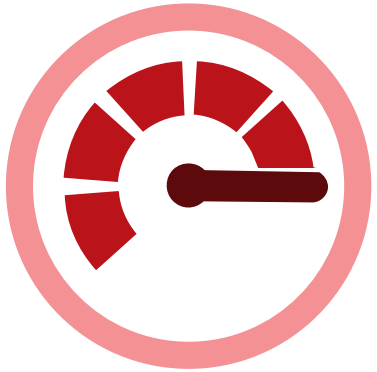

Database Recovery Advisor

SQL Sever Management Studio Database Recovery Advisor facilitates construction of restore plans that implement optimal correct restore sequences

Restore-plan algorithm: Improved for complex restore scenarios

Point-in-time restores: Simplified restoration of database to given point in time, and automatically includes backups relevant to desired restore point

Backup to Azure block blobs

Backup to Azure (SQL Server 2012)

Benefits:

- Near “bottomless” storage
- Off-site, geo-redundant
- No device management
- Remote accessibility

Limitations:

- Backup size up to 1 TB
- Restore speed

Backup to Azure block blobs (SQL Server 2016)

- 2x cheaper storage
- Backup striping and faster restore
- Maximum backup size is 12 TB+
- Granular access and unified credential story (SAS URIs)
- Supports all existing backup/restore features (except append)

Managed Backup

In SQL Server 2016, Managed Backup to Microsoft Azure uses new block blob storage for backup files

Stripe backup sets, enabling backup file sizes up to 12.8 TB

Other changes and enhancements to Managed Backup:

- Managed Backup used for system databases
- Support for databases in full, bulk logged, and simple recovery model
- Support for both automated and custom scheduling of backups
- Customized backup schedules – full backup and log backup

Backups and Stretch Databases

Backup of Stretch DB is a “shallow” backup only (backup/restore of local SQL Server hot data only)

Stretch DB feature ensures remote data is transactionally consistent with local data after each restore

Upon completion of local restore, SQL Server reconciles with remote using metadata

For more information:
https://msdn.microsoft.com/en-us/library/dn934993.aspx

Microsoft