Always On Failover Cluster Instances (SQL Server)

01/18/2017 • 9 minutes to read • Contributors

In this article

Benefits of a Failover Cluster Instance
Recommendations
Failover Cluster Instance Overview
Elements of a Failover Cluster Instance
SQL Server Failover Concepts and Tasks
Related Topics

THIS TOPIC APPLIES TO: SQL Server Azure SQL Database Azure SQL Data Warehouse Parallel Data Warehouse

As part of the SQL Server Always On offering, Always On Failover Cluster Instances leverages Windows Server Failover Clustering (WSFC) functionality to provide local high availability through redundancy at the server-instance level—a *failover cluster instance* (FCI). An FCI is a single instance of SQL Server that is installed across Windows Server Failover Clustering (WSFC) nodes and, possibly, across multiple subnets. On the network, an FCI appears to be an instance of SQL Server running on a single computer, but the FCI provides failover from one WSFC node to another if the current node becomes unavailable.

An FCI can leverage Availability Groups to provide remote disaster recovery at the database level. For more information, see Failover Clustering and Availability Groups (SQL Server).

> **Note**
>
> Windows Server 2016 Datacenter edition introduces support for Storage Spaces Direct (S2D). SQL Server Failover Cluster Instances support S2D for cluster storage resources. For more information, see Storage Spaces Direct in Windows Server 2016.

FEEDBACK

Failover Cluster Instances also support Clustered Shared Volumes (CSV). For more information, see [Understanding Cluster Shared Volumes in a Failover Cluster](#).

In this Topic:

- Benefits
- Recommendations
- Failover Cluster Instance Overview
- Elements of a Failover Cluster Instance
- SQL Server Failover Concepts and Tasks
- Related Topics

Benefits of a Failover Cluster Instance

When there is hardware or software failure of a server, the applications or clients connecting to the server will experience downtime. When a SQL Server instance is configured to be an FCI (instead of a standalone instance), the high availability of that SQL Server instance is protected by the presence of redundant nodes in the FCI. Only one of the nodes in the FCI owns the WSFC resource group at a time. In case of a failure (hardware failures, operating system failures, application or service failures), or a planned upgrade, the resource group ownership is moved to another WSFC node. This process is transparent to the client or application connecting to SQL Server and this minimize the downtime the application or clients experience during a failure. The following lists some key benefits that SQL Server failover cluster instances provide:

- Protection at the instance level through redundancy
- Automatic failover in the event of a failure (hardware failures, operating system failures, application or service failures)

  > **Important**
  >
  > In an availability group, automatic failover from an FCI to other nodes within the availability group is not supported. This means that FCIs and standalone nodes should not be coupled together within an availability group if automatic failover is an important component your high availability solution. However, this coupling can be made for your *disaster recovery* solution.

- Support for a broad array of storage solutions, including WSFC cluster disks (iSCSI, Fiber Channel, and so on) and server message block (SMB) file shares.

- Disaster recovery solution using a multi-subnet FCI or running an FCI-hosted database inside an availability group. With the new multi-subnet support in Microsoft SQL Server 2012, a multi-subnet FCI no longer requires a virtual LAN, increasing the manageability and security of a multi-subnet FCI.
- Zero reconfiguration of applications and clients during failovers
- Flexible failover policy for granular trigger events for automatic failovers
- Reliable failovers through periodic and detailed health detection using dedicated and persisted connections
- Configurability and predictability in failover time through indirect background checkpoints
- Throttled resource usage during failovers

Recommendations

In a production environment, we recommend that you use static IP addresses in conjunction the virtual IP address of a Failover Cluster Instance. We recommend against using DHCP in a production environment. In the event of down time, if the DHCP IP lease expires, extra time is required to re-register the new DHCP IP address associated with the DNS name.

Failover Cluster Instance Overview

An FCI runs in a WSFC resource group with one or more WSFC nodes. When the FCI starts up, one of the nodes assume ownership of the resource group and brings its SQL Server instance online. The resources owned by this node include:

- Network name
- IP address
- Shared disks
- SQL Server Database Engine service
- SQL Server Agent service
- SQL Server Analysis Services service, if installed
- One file share resource, if the FILESTREAM feature is installed

  At any time, only the resource group owner (and no other node in the FCI) is running its respective SQL Server services in the resource group. When a failover occurs, whether it be an automatic failover or a planned failover, the following sequence of events happen:

1. Unless a hardware or system failure occurs, all dirty pages in the buffer cache are written to disk.
2. All respective SQL Server services in the resource group are stopped on the active node.
3. The resource group ownership is transferred to another node in the FCI.
4. The new resource group owner starts its SQL Server services.
5. Client application connection requests are automatically directed to the new active node using the same virtual network name (VNN).

   The FCI is online as long as its underlying WSFC cluster is in good quorum health (the majority of the quorum WSFC nodes are available as automatic failover targets). When the WSFC cluster loses its quorum, whether due to hardware, software, network failure, or improper quorum configuration, the entire WSFC cluster, along with the FCI, is brought offline. Manual intervention is then required in this unplanned failover scenario to reestablish quorum in the remaining available nodes in order to bring the WSFC cluster and FCI back online. For more information, see WSFC Quorum Modes and Voting Configuration (SQL Server).

Predictable Failover Time

Depending on when your SQL Server instance last performed a checkpoint operation, there can be a substantial amount of dirty pages in the buffer cache. Consequently, failovers last as long as it takes to write the remaining dirty pages to disk, which can lead to long and unpredictable failover time. Beginning with Microsoft SQL Server 2012, the FCI can use indirect checkpoints to throttle the amount of dirty pages kept in the buffer cache. While this does consume additional resources under regular workload, it makes the failover time more predictable as well as more configurable. This is very useful when the service-level agreement in your organization specifies the recovery time objective (RTO) for your high availability solution. For more information on indirect checkpoints, see Indirect Checkpoints.

Reliable Health Monitoring and Flexible Failover Policy

After the FCI starts successfully, the WSFC service monitors both the health of the underlying WSFC cluster, as well as the health of the SQL Server instance. Beginning with Microsoft SQL Server 2012, the WSFC service uses a dedicated connection to poll the active SQL Server instance for detailed component diagnostics through a system stored procedure. The implication of this is three-fold:

- The dedicated connection to the SQL Server instance makes it possible to reliably poll for component diagnostics all the time, even when the FCI is under heavy load. This makes it possible to distinguish between a system that is under heavy load and a system that actually has failure conditions, thus preventing issues such as false failovers.
- The detailed component diagnostics makes it possible to configure a more flexible failover policy, whereby you can choose what failure conditions trigger failovers and which failure conditions do not.

- The detailed component diagnostics also enables better troubleshooting of automatic failovers retroactively. The diagnostic information is stored to log files, which are collocated with the SQL Server error logs. You can load them into the Log File Viewer to inspect the component states leading up to the failover occurrence in order to determine what cause that failover.

  For more information, see [Failover Policy for Failover Cluster Instances]

Elements of a Failover Cluster Instance

An FCI consists of a set of physical servers (nodes) that contain similar hardware configuration as well as identical software configuration that includes operating system version and patch level, and SQL Server version, patch level, components, and instance name. Identical software configuration is necessary to ensure that the FCI can be fully functional as it fails over between the nodes.

WSFC Resource Group
A SQL Server FCI runs in a WSFC resource group. Each node in the resource group maintains a synchronized copy of the configuration settings and check-pointed registry keys to ensure full functionality of the FCI after a failover, and only one of the nodes in the cluster owns the resource group at a time (the active node). The WSFC service manages the server cluster, quorum configuration, failover policy, and failover operations, as well as the VNN and virtual IP addresses for the FCI. In case of a failure (hardware failures, operating system failures, application or service failures) or a planned upgrade, the resource group ownership is moved to another node in the FCI.The number of nodes that are supported in a WSFC resource group depends on your SQL Server edition. Also, the same WSFC cluster can run multiple FCIs (multiple resource groups), depending on your hardware capacity, such as CPUs, memory, and number of disks.

SQL Server Binaries
The product binaries are installed locally on each node of the FCI, a process similar to SQL Server stand-alone installations. However, during startup, the services are not started automatically, but managed by WSFC.

Storage
Contrary to the availability group, an FCI must use shared storage between all nodes of the FCI for database and log storage. The shared storage can be in the form of WSFC cluster disks, disks on a SAN, Storage Spaces Direct (S2D), or file shares on an SMB. This way, all nodes in the FCI have the same view of instance data whenever a failover occurs. This does mean, however, that the shared storage has the potential of being the single point of failure, and FCI depends on the underlying storage solution to ensure data protection.

Network Name
The VNN for the FCI provides a unified connection point for the FCI. This allows applications to connect to the VNN without the need to know the currently active node. When a failover occurs, the VNN is registered to the new active node after it starts. This process is transparent to the client or application

connecting to SQL Server and this minimize the downtime the application or clients experience during a failure.

Virtual IPs
In the case of a multi-subnet FCI, a virtual IP address is assigned to each subnet in the FCI. During a failover, the VNN on the DNS server is updated to point to the virtual IP address for the respective subnet. Applications and clients can then connect to the FCI using the same VNN after a multi-subnet failover.

SQL Server Failover Concepts and Tasks

Concepts and Tasks	Topic
Describes the failure detection mechanism and the flexible failover policy.	Failover Policy for Failover Cluster Instances
Describes concepts in FCI administration and maintenance.	Failover Cluster Instance Administration and Maintenance
Describes multi-subnet configuration and concepts	SQL Server Multi-Subnet Clustering (SQL Server)

Related Topics

Topic descriptions	Topic
Describes how to install a new SQL Server FCI.	Create a New SQL Server Failover Cluster (Setup)
Describes how to upgrade to a SQL Server 2017 failover cluster.	Upgrade a SQL Server Failover Cluster Instance
Describes Windows Failover Clustering Concepts and provides links to tasks related to Windows Failover Clustering	Windows Server 2008: Overview of Failover Clusters Windows Server 2008 R2: Overview of Failover Clusters
Describes the distinctions in concepts between nodes in an FCI and replicas within an availability group and considerations for using an FCI to host a replica for an availability group.	Failover Clustering and Availability Groups (SQL Server)

Note

The feedback system for this content will be changing soon. Old comments will not be carried over. If content within a comment thread is important to you, please save a copy. For more information on

the upcoming change, we invite you to read our blog post.

0 Comments

Sign in

26 people listening

+ Follow

Post comment as...