Hybrid Cloud Scale-Out Design for Microsoft SQL Server 2014 with Cisco Intercloud Fabric

December 2015

Contents

Introduction

Industry trends indicate a growing movement among organizations to hybrid cloud designs. Organizations are choosing these designs to acquire additional on-demand computing and storage resources, thereby eliminating the need to build for peak capacity within their own data centers. Public clouds do not require the initial capital investments necessary to build out a company's own private cloud. In addition, a public cloud can better absorb a company's need for elasticity by providing almost unlimited pay-as-you-grow expansion. However, although hybrid cloud models are conceptually and financially very attractive, customers are often reluctant to place their applications in the public cloud, away from their own premises. When an organization deploys an application or part of an application in the public cloud, it wants to be sure that the transition from the private cloud to a hybrid model is not only operationally feasible, but also that the company retains data access and control in the new architecture.

Cisco Intercloud Fabric™ facilitates secure hybrid cloud creation, helping enable customers to scale applications securely and conveniently. Using Cisco Intercloud Fabric, customers can access additional computing power while extending their security and other control policies to the public cloud. Cisco Intercloud Fabric also allows customers to place workloads across heterogeneous environments in multiple provider clouds.

Main Benefits

Cisco Intercloud Fabric offers these main benefits:

- A single point of management and control for virtual workloads across multiple provider clouds
- A choice of cloud providers, such as Amazon Web Services (AWS) and Microsoft Azure, or multiple Intercloud fabric provider–based clouds
- Highly secure and scalable connectivity to extend private clouds to service provider clouds
- Enforcement of consistent network and workload policies throughout the hybrid cloud
- Workload mobility to and from service provider clouds for virtual workloads

Audience

The audience for this document includes sales engineers, field consultants, professional services, IT managers, partner engineers, and customers who want to take advantage of an infrastructure that is built to scale on demand while maintaining consistent security and control policies.

Scope

This document describes the architecture and design details for deploying and scaling Microsoft SQL Server 2014 using Cisco Intercloud Fabric for Business. With Cisco Intercloud Fabric for Business, you can extend your data center or private cloud to the public cloud, allowing you to acquire the additional capacity you need with no border between your internal cloud and the external cloud. This document discusses a reference design using SQL Server 2014 in a VMware-based internal data center and describes how to scale this application by using resources in Microsoft Azure.

Technology Use Case

You can use the solution described in this document to achieve cloud benefits without moving mission-critical on-premises databases. You can incorporate the cloud for backup, reporting, and disaster recovery to take advantage of the cloud without the risk entailed in moving your mission-critical on-premises databases.

Cisco Intercloud Fabric

Cisco Intercloud Fabric architecture is available in two configurations to address enterprise and service provider customers:

- Cisco Intercloud Fabric for Business
- Cisco Intercloud Fabric for Providers

Cisco Intercloud Fabric for Business is intended for enterprise customers who want to transparently extend their private clouds to provider cloud environments. Cisco Intercloud Fabric for Providers is intended for provider-managed cloud environments, to allow the service provider's enterprise customers to transparently extend their private cloud environments into the provider's cloud.

This document focuses on an enterprise customer deployment environment and therefore discusses only Cisco Intercloud Fabric for Business.

Cisco Intercloud Fabric Architecture for Microsoft SQL Server 2014

Cisco Intercloud Fabric for Business is intended for enterprise customers who want to transparently extend their private cloud into public cloud environments while keeping the same level of security and policy across environments.

Figure 1 shows the connectivity between the private cloud (enterprise data center) and Microsoft Azure cloud (provider cloud) in the design used for this study.

Figure 1.　　Cisco Intercloud Fabric Architecture for Microsoft SQL Server 2014
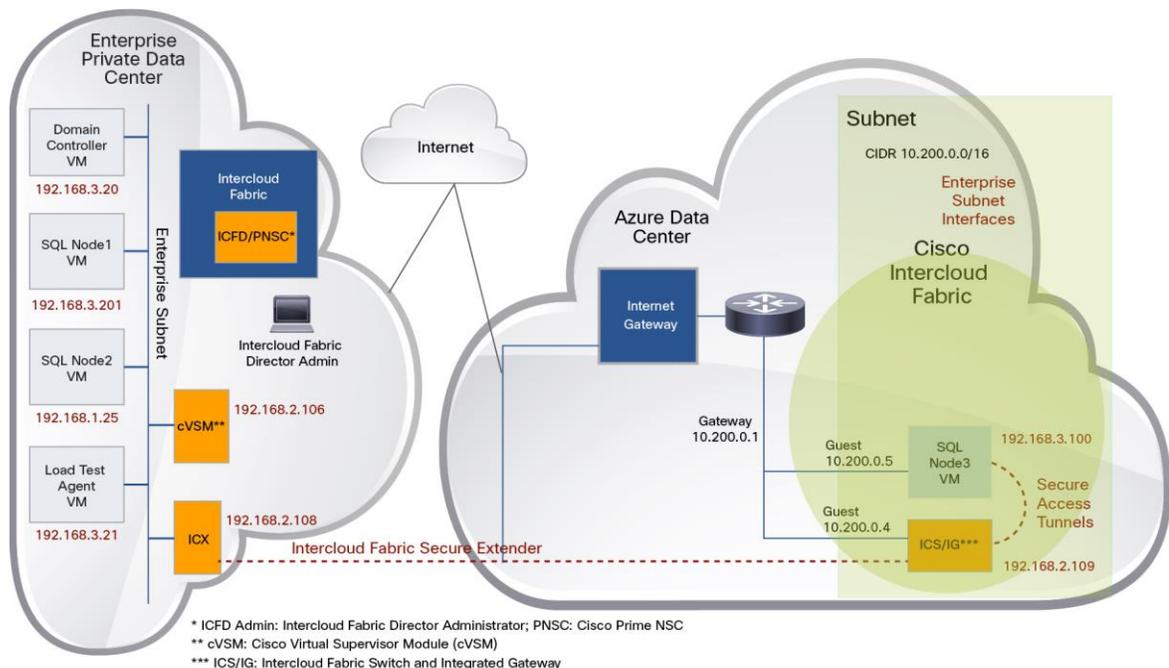
The Intercloud fabric director, Cisco Virtual Supervisor Module (cVSM), and Intercloud fabric secure extender infrastructure components are deployed in the private cloud (enterprise data center), and the Intercloud Fabric Switch (ICS) and integrated gateway deployed in the public cloud (Microsoft Azure). For additional information, refer to [Cisco Intercloud Fabric Configuration Guide, Release 2.3.1](#).

Deployment of the cloud services router integrated gateway (Cisco Intercloud Fabric integrated router) allows routing of the overlay extended networks within the cloud provider environment. The Cisco Intercloud Fabric router was configured to allow external clients (that is, clients not connected to the enterprise) to access the SQL reporting services using the public Internet.

All data in motion (SQL replication data) is cryptographically isolated and encrypted within the Intercloud fabric secure extender. This data includes traffic exchanged between the private and public clouds (site to site) and the virtual machines running in the cloud (virtual machine to virtual machine). A Datagram Transport Layer Security (DTLS) tunnel is created between endpoints to more securely transmit this data. DTLS is a User Datagram Protocol (UDP)–based, highly secure transmission protocol.

The SQL Server primary database instance is deployed in the private cloud (enterprise data center), and secondary instances are deployed in the Azure cloud. The replication between the primary and secondary databases occurs through this secure tunnel. This process requires the opening of appropriate ports in the enterprise firewall: for example, in this setup, ports 443 and 22 are opened.

Microsoft SQL Server 2014 AlwaysOn in Hybrid Cloud

SQL Server 2014 AlwaysOn availability groups can protect multiple databases by replicating all database transactions on up to eight replica servers. Combining infrastructure as a service (IaaS) with AlwaysOn availability groups enables customers to create and host replicas in the hybrid cloud, creating a high-availability solution without the need to set up their own high-availability and disaster-recovery site. If a primary site or primary SQL Server database fails, the database instance in the hybrid cloud can provide data until the primary on-premises database is available. Customers can gain these high-availability benefits without the need for any additional capital expenditures. With the hybrid cloud, you pay only for the resources that you require.

For the purpose of this study, a primary server and a secondary replica are configured with synchronous commit on premises, and a third replica is configured with asynchronous replication in the Azure cloud. The Intercloud Fabric Switch (ICS) can perform the routing and Network Address Translation (NAT) functions using the integrated gateway feature with the Azure cloud.

Organizations have complete control over both the operating system and the installed applications and can achieve almost complete compatibility with the on-premises SQL Server and take advantage of AlwaysOn availability groups. They can use secondary replicas for high availability, disaster recovery, reporting, and backup Figure 2.

Figure 2. Microsoft SQL Servers in a Microsoft Windows Failover Cluster with the AlwaysOn Feature Enabled in Hybrid Cloud
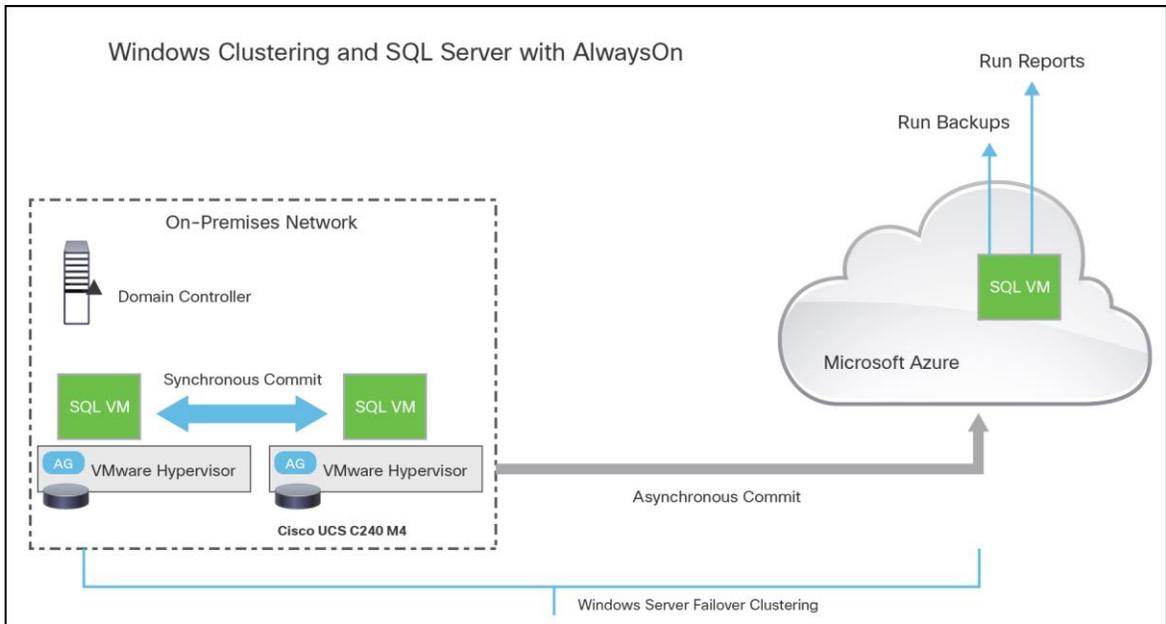
Hardware and Software Versions

Table 1 lists the hardware and software versions used in this deployment model.

Table 1. Hardware and Software Versions

Layer	Device	Image	Comments
Computing	Cisco UCS® C240 M3 Rack Server	Version 2.0 (6d)	Cisco Unified Computing System™ (Cisco UCS) server
Network adapter	Cisco UCS Virtual Interface Card (VIC) 1225	Version 4.0 (5b)	Cisco® VIC
Cisco Intercloud Fabric software	Cisco Intercloud Fabric	Version 2.3.1	Cisco Intercloud Fabric Software
VMware software	VMware ESXi 5.5.0	Version 2068190	VMware ESXi
Microsoft software	Microsoft Windows 2012 R2	Version 2012 R2	Data center operating system
Microsoft application	Microsoft SQL Server 2014	Version 2014	Enterprise Edition
HammerDB open-source software	HammerDB	Version 2.17	Load-test tool

Microsoft SQL Server 2014 AlwaysOn Architecture

This section provides an overview of the end-to-end design of the reference architecture for SQL Server 2014 with the AlwaysOn feature implemented in Cisco Intercloud Fabric.

Figure 3 shows Microsoft Windows Server Failover Clustering (WSFC) with three SQL Server 2014 virtual machines and with the AlwaysOn availability group feature enabled. This feature combination provides infrastructure support for the high-availability and disaster-recovery features of the hosted database server.

Figure 3.   Microsoft SQL Servers with Cisco Intercloud Fabric in a Hybrid Cloud
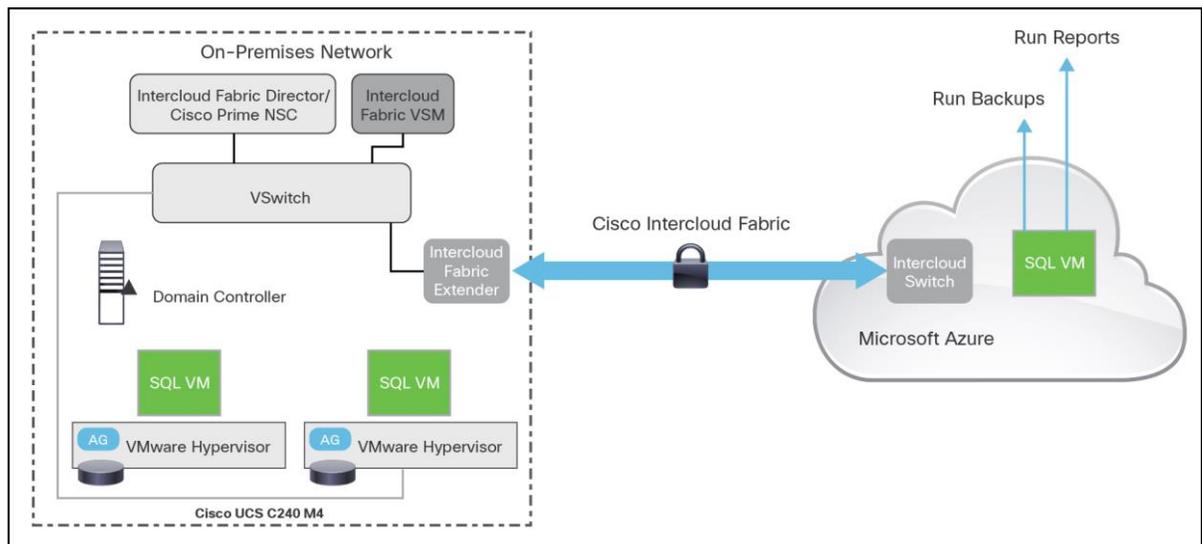
An AlwaysOn availability group is composed of replicas of one or more databases. You can have one primary replica and one to eight secondary replicas. The sample configuration in this document shows multiple secondary replicas. A secondary replica is a copy of the database from the primary replica. When the database is modified, changes are replicated to all secondary replicas.

This sample configuration also uses an availability group listener to direct the incoming connections to the primary replica or to a read-only secondary replica. The listener provides fast application failover after an availability group fails. An availability group listener is a virtual network name (VNN) to which the clients can connect to access a database in a primary or a secondary replica of an AlwaysOn availability group.

The VNN is registered in the Domain Name System (DNS) and is always owned by the SQL Server instance on which the primary replica resides. All the IP addresses that are supplied while configuring the availability group listener are registered in DNS under the same VNN. In this document, the availability group listener is referred to as SQL AG.

After SQL AG is created, verify that the clients can connect. The only change in the application connection is that instead of pointing to a specific server in the connection string, you point to SQL AG. SQL AG can connect only using TCP, and it can be resolved by your local DNS to the list of IP addresses and the TCP port (1433) mapped to the VNN. When failover occurs, client connections are reset, and the ownership of the availability group listener (SQL AG) moves to the SQL Server instance that takes over the primary replica role. The VNN endpoint is then bonded to the new IP addresses and TCP port of the new primary replica instance. Depending on the client,

automatic reconnection to the availability group will occur, or you may have to manually restart the affected application or connection.

In this document, the availability group is configured with both synchronous and asynchronous commit modes.

Figure 4 shows three virtual machine network interface cards (NICs)—host management, Windows cluster, and SQL availability group listener—created on SQL nodes 1, 2, and 3 with different unique VLANs that separate the traffic.
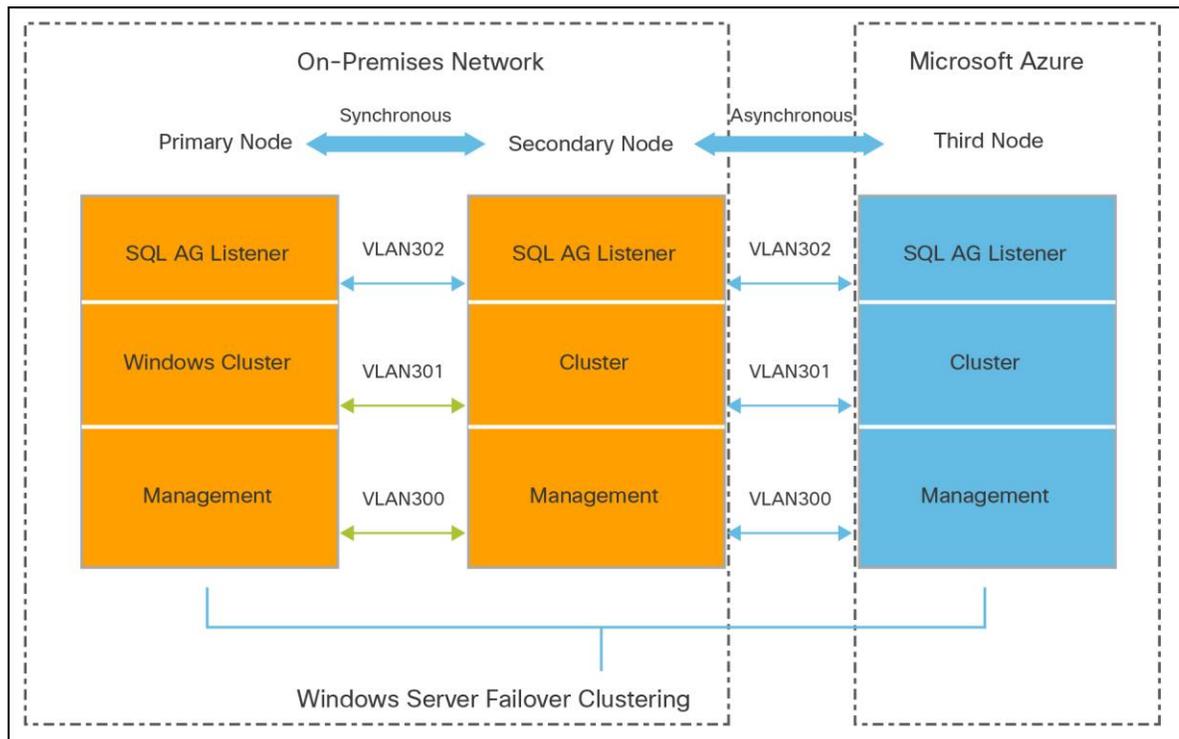
Figure 4.   Virtual Machine Network Interfaces

Table 2 and Table 3, respectively, show the configuration details for the SQL Server 2014 AlwaysOn availability group and availability replica group features.

Table 2.   Microsoft SQL Server AlwaysOn Availability Group

Microsoft SQL Server Availability Group	Member (Instance)	Default Role	Database	Database Configuration	Database Size	Memory Settings
Availability group database	SQL2014-NODE1\SQLDBN1 SQL2014-NODE2\SQLN2 SQL2014-NODE2\SQLN3	Primary Secondary Secondary	Online transaction processing (OLTP) workloads	Database and log are on the same drive	1 GB	Minimum: 8 GB Maximum: 16 GB

Table 3. Microsoft SQL Server AlwaysOn Availability Replica Group

Microsoft SQL Server Availability Group	Service Instance	Automatic Failover	Synchronous Commit	Allow Readable Secondary	Replication Network
Availability group database	SQL2014-NODE1\SQLDBN1 SQL2014-NODE2\SQLN2 SQL2014-NODE2\SQLN3	Yes Yes No	Yes Yes No	Yes Yes	SQL services network

Verification Summary

Testing verified that failover occurred successfully. Additional tests were performed to verify that the schema and application structures of the primary replica on the primary node and the secondary readable replica on the secondary node are consistent after failover.

Conclusion

Cisco Intercloud Fabric for Business facilitates the creation of a secure hybrid cloud, providing customers the capability to scale up their applications securely and conveniently. Using Cisco Intercloud Fabric, customers can access additional computing power while maintaining security and control of their applications' access to on-premises authentication and database servers. The Intercloud fabric director also allows you to upload the virtual machine templates to the cloud for instant application server deployment and provides the capability to migrate customer workloads between the private and public clouds.

Cisco Intercloud Fabric has proven to be an easy-to-manage, easy-to-deploy, self-contained infrastructure for an enterprise-class Microsoft SQL Server 2014 database server. This framework has been proven to yield high availability, efficiency, and lower infrastructure cost.

For More Information

Cisco Intercloud Fabric Release 2.3.1 release notes

- Cisco Intercloud Fabric Release Notes, Release 2.3.1

Cisco Intercloud Fabric for Business:

- Cisco Intercloud Fabric Configuration Guide, Release 2.3.1

Cisco Intercloud Fabric: Hybrid cloud with choice, consistency, control, and compliance:

- http://www.cisco.com/c/en/us/td/docs/solutions/Hybrid_Cloud/Intercloud/Intercloud_Fabric.html