Topologies for SharePoint 2013

Microsoft® SharePoint®

TOPOLOGY DESIGN PRINCIPLES

Overview

The traditional three-tier roles of a Microsoft® SharePoint® 2013 farm can be deployed on a single server for evaluation or development, or on many servers. The three-tier roles include:

- Web server role — Fast, light-weight server which responds to user requests for web pages. All web servers in a farm are mirrors of each other and are load balanced.
- Application server role — Provides the service features of SharePoint products and technologies. An application server often provides all or a subset of service features. Multiple redundant application servers can be load balanced.
- Database server role — Stores content and service data. All databases can be assigned to one database server. or databases can be spread across multiple servers. All databases can be clustered or mirrored for failover protection.

In a small farm, server roles can be combined on one or two servers. For example, web server and application server roles can be combined on a single server or on two or more servers to achieve redundancy.

Service applications

Service applications are services that are shared across sites within a farm (for example, Search and Excel Services). Some service applications can be shared across multiple farms. Some services support partitioning.

Service applications are deployed to the application server tier. Some services include multiple components, and deployment of these components requires planning. For example:

- The Search service application includes multiple application components and multiple databases.
- The User Profile service application includes multiple databases.

Each service application is associated with at least one service on the Services on Server page in Central Administration.

Services on server

The Services on Server page in Central Administration lists services that are started or stopped on specific servers in the farm:

- Some of these services are associated with service applications. After you deploy service applications to the farm, go to the Services on Server page and ensure that the associated services are started on the appropriate servers.
- Some of these services are not associated with service applications.

After you plan the farm topology, see Plan services on server in the TechNet library to plan the mapping of services to server applications.

Note: To deploy search components to servers, you use the Search service application pages in Central Administration instead of the Services on Server page.

Virtual topologies

This model provides examples of virtualized topologies. Virtualized topologies depend on the capacity of physical hosts, desired ratio or virtual machines to hosts, and the underlying virtualization technology.

Server roles

Web server

- Hosts web pages, Web services, and Web Parts that are necessary to process requests served by the farm.
- Directs requests to the appropriate application servers.
- In dedicated services farms, this role is not necessary because web servers at remote farms contact application servers directly.

Application server roles

Use the Services on Server page in Central Administration to assign services to specific application servers.

- In many farms, all services will run on two identically configured application servers for redundancy.
- The Search service application automatically configures the necessary services on application servers. Using the Services on Server page is not necessary.
- After deployment, look for services that consume a disproportionate amount of resources and consider placing these services on dedicated hardware.

Database server

In a small farm environment, all databases can be deployed to a single server. In larger environments, group databases by roles and deploy these to multiple database servers.

TOPOLOGY CONCEPTS FOR SHAREPOINT 2013

Smallest fault-tolerant farm

The smallest fully redundant physical farm incorporates six servers — two for each tier. User requests are automatically load-balanced across the web servers and application servers are utilized equally.

For SharePoint 2013, the query processing component replaces the query role of previous versions. The query processing component requires more resources and is not recommended for web servers unless these are sized appropriately.

Use SQL Server clustering, mirroring, or AlwaysOn for the database servers. AlwaysOn requires SQL Server 2012.

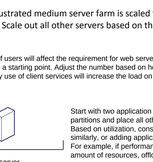

Important: The query processing component in SharePoint 2013 offloads much of the CPU and disk load from SQL Server. The footprint and performance requirements for SQL Server in SharePoint 2013 are lower than the previous product version. As a result of this architecture improvement, the query processing component requires more local resources than previous versions. The query role can be combined with the Web server role on a server only if there are enough resources. Running both of these roles on a single virtual machine requires a 6-8-core VM and a physical host that runs Windows Server 2012. A 4-core VM does not provide enough resources for both the query processing component and the Web server role.

Search optimized farm

A search-optimized farm separates the query processing component and index component to dedicated application servers. The remainder of the search components and all other application roles remain on two all-purpose application servers.

If crawling is producing more traffic on web servers than user requests, you can dedicate one or more web servers for crawling. We recommend this in environments that crawl large amounts of data. In SharePoint 2013, it is not necessary or recommended to configure affinity for these servers on the load balancer.

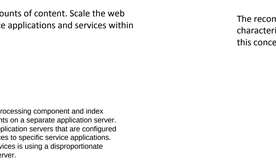

Distributed cache

The distributed cache feature is enabled by default and the Distributed Cache service is automatically started on all web and application servers in a farm. Distributed cache improves performance by:

- Caching social data, such as news feeds.
- Caching authentication tokens.

In very large environments distributed cache can be offloaded to dedicated servers.

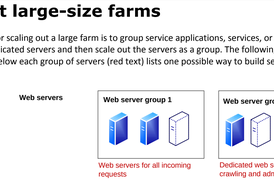

Stretched farms are not supported for SharePoint 2013: Redundancy and failover between closely located data centers that are configured as a single farm ("stretched" farm) are not supported in SharePoint 2013. All servers that belong to a server farm must be physically located in the same datacenter.

Office Web Apps Server

Office Web Apps Server is a separate server product that can:

- Serve multiple SharePoint Server farms for viewing and editing.
- View files from Exchange Server, Microsoft Lync.
- Integrate with URL-accessible file servers.

By separating Office Web Apps from the SharePoint farm, you can update servers more frequently and manage scale and performance independent of the SharePoint environment. Office Web Apps Server can be used with all versions of SharePoint 2013. The Office Web Apps Server architecture does not include a database.

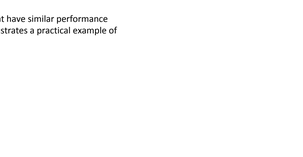

Request management and load balancing

Request Management is a feature that gives SharePoint farms control over incoming requests and how these are routed. Routing rules are prioritized and apply logic to determine the nature of requests and to apply the most appropriate response, such as the following types of actions:

- Route requests to web servers that have good health characteristics.
- Identify and block known bad requests .
- Route requests of specific types (such as search) to specific servers in the farm.

Request management does not replace the role of a load balancer and it is not enabled by default.

The Request Management component can run in integrated mode on chosen web servers in a farm. Alternatively, the Request Management component can run on dedicated servers that are not part of the SharePoint farm.

- Integrated mode — Request Management runs on the web servers you choose in a farm. This mode is appropriate for most environments (shown below).
- Dedicated mode — Servers in a separate Request Management farm sit between the hardware load balancer and one or more SharePoint farms. This mode is appropriate for large-scale environments. With this configuration, Request Management can serve several SharePoint farms. A Request Management farm can be scaled independently based on utilization (not shown).

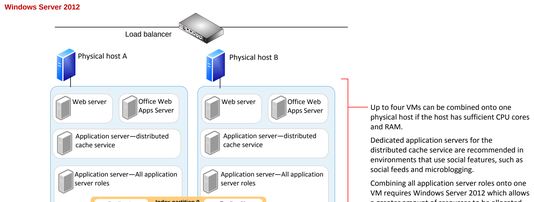

Request Management does not apply to Office Web Apps Server.

Request Manager runs within SharePoint under SPRequestModule and responds to HTTP requests only.

To configure Request Manager:

- Start the Request Management service on desired web servers.
- Configure rules by using Windows PowerShell.

EXAMPLE TOPOLOGIES

Limited deployments (1-2 servers)

Limited deployments are typically used for product evaluation, development and testing, or for environments that have limited numbers of users and don't require fault-tolerance.

One-server farm

Evaluation or <100 users

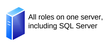

Two-tier farm

Up to 10,000 users

Development environments

Mimic a three-tier environment by using virtualization.

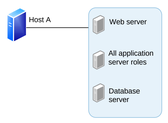

Small multipurpose farms (3-4 servers)

Small farm architectures serve a larger number of users and scale out based on how heavily services are used. Not all small farms are fault-tolerant.

Three-server virtualized farm

Use virtualization to maximize the potential of a smaller number of servers

Two web servers are predicted to serve 10,000-20,000 users.

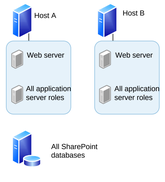

High availability — For environments above 1,000 users, two servers utilizing SQL Server clustering, mirroring, or AlwaysOn is recommended. AlwaysOn requires SQL Server 2012.

Four-server physical farms

Add a dedicated application server for environments with moderate service usage.

If fault-tolerance of services is more important than user performance, configure the farm with redundant application servers instead of redundant Web servers.

Smallest fault-tolerant farm utilizing virtualization

All farm server roles virtualized and distributed across two or four host servers (depending on the operating system) to provide fault tolerance using the minimum number of servers.

Windows Server 2008 R2

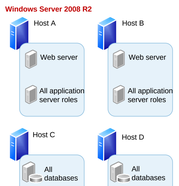

SQL Server installed and configured to support SQL clustering, mirroring, or AlwaysOn. AlwaysOn requires SQL Server 2012.

Windows Server 2012

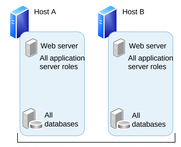

SQL Server installed and configured to support SQL clustering, mirroring, or AlwaysOn. AlwaysOn requires SQL Server 2012.

Virtualizing with Windows Server 2008 R2 versus Windows Server 2012 — Windows Server 2012 allows a greater number of resources to be allocated to each virtual machine, and thus allows running a SharePoint farm with fewer virtual machines than Windows Server 2008 R2. Due to the virtual architectures that are possible, this topology model indicates which operating system a virtual topology is designed for.

Legend for database icons

Medium farm architectures (6+ servers)

Medium farm architectures can be multi-purpose or optimized for specific purposes. Medium-size farms are fully fault-tolerant. Some environments might require more web servers. Factor 10,000 users per web server as a starting point.

Six-server physical farm

Six-server virtualized farm

Windows Server 2008 R2

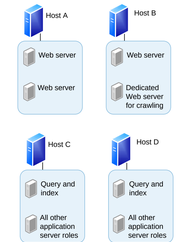

Eight-server physical farm optimized for search

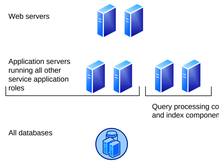

Eight-server virtualized farm

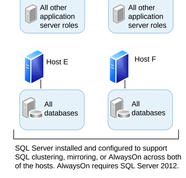

SQL Server installed and configured to support SQL clustering, mirroring, or AlwaysOn across both of the hosts. AlwaysOn requires SQL Server 2012.

Medium farm with Office Web Apps Server and articulated search roles

A medium-size, fault-tolerant, virtual environment that includes Office Web Apps Server VMs and a plan for scaling out databases. This diagram also calls out all search application roles.

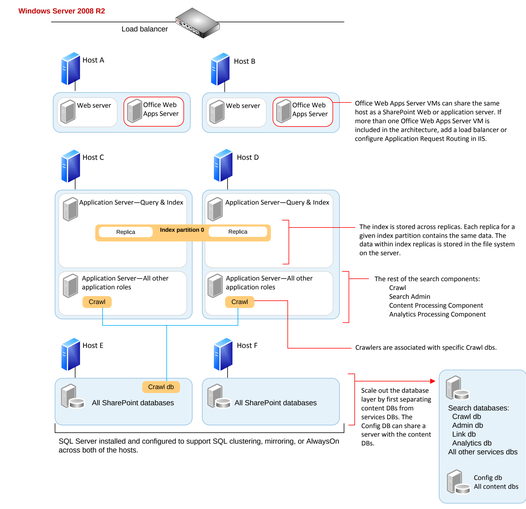

SCALING FARMS WITH SERVER GROUPS AND STORAGE GROUPS

Server groups

The recommendation for scaling out a farm is to group services or databases that have similar performance characteristics onto dedicated servers and then scale out the servers as a group.

These topology examples group service applications and related components (for example, databases) into several different logical groupings that can be used as a starting point. In large environments, the specific groups that evolve for a farm depend on the specific demands for each service.

Note: *Server groups* is a planning concept. This term and concept is not found in Central Administration.

Scaling out search

As illustrated, the farm on the right is designed to crawl 10 million items (~800 GB of data). The correlation between items and volume of data will vary depending on the types of data that are crawled. It is important to understand the characteristics of the data within the environment. Above 40 million items, consider a dedicated search farm. The following table provides starting-point numbers for search components based on number of items.

For more information and examples on scaling search farms, see the following model: Example Search Architectures for Enterprise Search.

Scaling out medium-size farms

Scale out medium-size farms based on the workload. This illustrated medium server farm is scaled to host sites with large amounts of content. Scale the web servers and content database servers as needed for growth. Scale out all other servers based on the utilization of other service applications and services within the farm and the volume of content the farm will host.

Starting-point numbers for search components based on number of items *

Number of items	Index Components and Partitions	Query Processing Components	Content Processing Components	Analytics Processing Components	Crawlers	Crawl databases	Link database	Analytics reporting database	Search Administration Component
General guidance	Add 1 index partition per 10 million items	Use 2 query processing components for redundancy. Above 80 million items, increase to 4.				Add 1 crawl database per 20 million items.	Add 1 link database per 60 million items.	Add one analytics reporting database for each 500K unique items viewed each day or every 10-20M total items.	Use 2 search administration components for redundancy, for all farm sizes.
10 million	2 components 1 partition	2	2	2	2	1	1	variable	2
10- 40 million	8 components 4 partitions	2	4	2	2	2	1	variable	2
100 million	20 components 10 partitions	4	6	6	2	5	2	variable	2

* This guidance is intended for enterprise search within an organization. Guidance for Internet sites is provided in a different document.

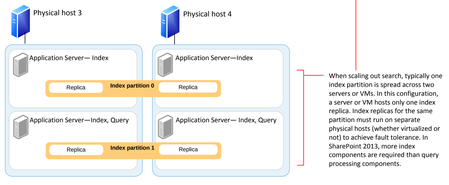

Scaling out large-size farms

The recommendation for scaling out a large farm is to group service applications, services, or databases that have similar performance characteristics onto dedicated servers and then scale out the servers as a group. The following topology illustrates a practical example of this concept. The text below each group of servers (red text) lists one possible way to build server groups.

Storage groups

Storage groups is a concept in which similar types of databases are grouped together and scaled out independent of the rest of the databases based on need. All databases within a storage group are treated the same with backup procedures and restore protocols. The best practice is to include the configuration database with the content database group.

Medium virtual farm consolidated onto 4 physical hosts

If the cost of server hardware is a limiting factor (as opposed to the cost of licensing), four VMs can be consolidated onto a single physical host given enough server resources.

Windows Server 2012

Microsoft®