SharePoint 2013 Logical Architecture

Table Of Contents

Plan logical architectures for SharePoint 2013

SharePoint 2013

Applies to: *SharePoint Server 2013, SharePoint Foundation 2013*

Topic Last Modified: *2013-12-18*

Summary: Learn how to visualize the architecture of a SharePoint 2013 deployment by using the logical architecture articles and other resources in this section.

The following downloadable resources, articles on TechNet, video recordings, and related resources provide information about logical architecture planning.

TechNet articles about logical architecture planning

The following articles about logical architecture planning are available to view online. Writers update articles on a continuing basis as new information becomes available and as users provide feedback.

	Content	Description
	SharePoint 2013 design samples: Corporate portal and extranet sites	Use the design samples to see a generic corporate deployment of SharePoint 2013.
	Plan service deployment in SharePoint 2013	Lists the services on the Services on Server page in the SharePoint Central Administration website and provides deployment guidance for these services. Many of these services are associated with service applications.
	Plan logical architectures for SharePoint 2013	Learn how to plan and implement host-named site collections in SharePoint 2013 and how path-bases site collections might affect your environment.
	Plan alternate access mappings for SharePoint 2013	Learn how to plan for alternate access mappings in SharePoint 2013.
	Configure alternate access mappings for SharePoint 2013	Learn how to configure alternate access mappings in SharePoint 2013.
	Update a web application URL and IIS bindings for SharePoint 2013	Learn how to update a web application and IIS bindings for SharePoint 2013.

Additional resources about logical architecture planning

The following resources about logical architecture planning are available from other subject matter experts.

	Content	Description
Microsoft | TechNet	• Resource Center: Architecture design for SharePoint 2013 IT pros	Visit the TechCenter to access videos, community sites, documentation, and more.

Plan service deployment in SharePoint 2013

SharePoint 2013

Applies to: *SharePoint Server 2013, SharePoint Foundation 2013*

Topic Last Modified: *2013-12-18*

Summary: Learn about the services on the Services on Server page in the SharePoint Central Administration website and deployment guidance for these services. Many of these services are associated with service applications.

In this article:

About 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 deploying service applications to the farm, go to the Services on Server page and make sure that the associated services are started on the appropriate servers. Some of these services are not associated with service applications.

Note:
Note: Search components are deployed by using Windows PowerShell instead of the **Services on Server** page.

After planning the farm topology, use this article to make sure that the appropriate services on server are started for each server in the farm. This may not be necessary for small farms in which all of the service applications and services are run on one server or two redundant servers.

This article lists recommended configuration of services and service applications for two different architecture approaches:

- **Streamlined topologies** — The distribution of services and other components in a farm is intended to maximize system resources of server hardware. Streamlined architectures include front-end servers, batch-processing servers, and database servers. You can use the Services on Server Install Worksheet for Streamlined Topologies to track decisions about which services are started on each server.
- **Traditional topologies** — Topologies are based on traditional approaches to building architectures with Web servers, application servers, and database servers. You can use the Services on Server Install Worksheet for Traditional Topologies to track decisions about which services are started on each server.

For more information about these topology approaches, see the following models:

- Streamlined Topologies for SharePoint Server 2013
- Traditional Topologies for SharePoint Server 2013

Services on server for streamlined topologies

Table: Service applications and Services on Server for streamlined topologies

Service application	Services on server	Server recommendation	Additional information
Access Services	Access Services	Front-end server	Enables users of the Access 2013 client to create new Access service applications. View, edit, and interact with Access Services databases in a browser.
Access Services 2010	Access Database Service 2010	Front-end server	Allows continued maintenance of SharePoint 2010 Access service applications by using Access 2010 clients and Access 2013 clients. Does not allow users to create new applications.
App Management Service	App Management Service	Front-end server	
Business Data Connectivity	Business Data Connectivity service	Front-end server	
Excel Services	Excel Calculation Services	Front-end server or specialized server	
Machine	Machine	Batch-processing server	The Machine Translation Service uses SharePoint OAuth for file management. The Machine

Translation Service	Translation Service		Translation Service Application's app pool account needs permissions to the User Profile Application to function correctly. The Machine Translation Service connects to the Microsoft Translator online service. Microsoft Translator may collect some data to improve the quality of translations.
Managed Metadata Service	Managed Metadata Web Service	Front-end server	
Microsoft SharePoint Foundation Subscription Settings Service	Microsoft SharePoint Foundation Subscription Settings Service	Front-end server	You must use Windows PowerShell to deploy this service. Start this service if you have deployed service applications in multi-tenant mode or if the farm includes sites that use site subscriptions. This service stores settings and configuration data for tenants in a multi-tenant environment. After you start this service, web applications consume the service automatically.
PerformancePoint	PerformancePoint Service	Batch-processing server or specialized server	
PowerPoint Conversion	PowerPoint Conversion Service	Batch-processing server	This service converts PowerPoint presentations to other formats. It typically runs on one or more application servers. It starts one or more worker processes to perform conversions. When actively converting, a worker process may use up a complete processor core. Memory usage depends on the size and content of files being converted. You can use Windows PowerShell cmdlets to control the number of worker processes that are used. Several other configuration options are also available through Windows PowerShell cmdlets.
Search	Lotus Notes Connector	Batch-processing server or specialized servers for search	Refer to TechNet documentation to learn how to configure Lotus Notes Connector which crawls data on a Lotus Domino server.
Search	Search Host Controller Service	Servers that run search components	This service manages the search topology components. The service is automatically started on all servers that run search topology components.
Search	Search Query and Site Settings Service	Servers that run the query processing component	This service load balances queries within the search topology. It also detects farm-level changes to the search service and puts these in the Search Admin database. The service is automatically started on all servers that run the query processing component.
Search	SharePoint Server Search	Servers that run search components	This service crawls content for the search index. This service is automatically started on all servers that run search topology components. The service cannot be stopped or started from the Services on Server page.
Secure Store Service	Secure Store Service	Front-end server	
Usage and Health Data Collection			This service application has no associated service on server.
User Profile	User Profile Service	Front-end server	
User Profile	User Profile Synchronization Service	Batch-processing server	
Visio Graphics Service	Visio Graphics Service	Front-end server	
Word Automation Service	Word Automation Services	Batch-processing server	Performs automated bulk document conversions. When actively converting, this service will fully use one CPU for each worker process (configured in Central Administration). If the service is started on multiple servers, a job will be shared across all the servers.
Work Management	Work Management Service	Batch-processing server	
	Central Administration	Front-end server	This service runs the Central Administration site. It can be put on a batch-processing server if security policies of an organization mandate this.
	Claims to Windows Token Service	Servers that run Excel Calculation Services and PerformancePoint as well as any other service applications that rely on this service to delegate a SharePoint user identity to an external data source	
	Distributed Cache	Front-end server OR Distributed Cache and Request Management tier (not both)	By default this service is started on all web servers and application servers in a farm.

	Document Conversions Launcher Service	Batch-processing servers	Schedules and starts the document conversions on a server.
	Document Conversions Load Balancer Service	Batch-processing servers	Balances document conversion requests from across the server farm. Each web application can have only one load balancer registered with it at a time.
	Microsoft SharePoint Foundation Incoming E-Mail	Front-end servers	Typically, this service runs on a web server. If you have to isolate this service, you can start it on an application server.
	Microsoft SharePoint Foundation Sandboxed Code Service	Front-end servers	Start this service on computers that run sandboxed code in the farm. This can include front-end servers and batch-processing servers. This service runs code that is deployed as part of a sandboxed solution in a remote, rights-restricted process and measures the server resources that are used during execution against a site collection-scoped, daily quota.
	Microsoft SharePoint Foundation Web Application	Front-end servers, batch-processing servers, plus the Distributed Cache and Request Management servers if these servers are implemented	This service provides web server functionality. It is started by default on web servers. Custom features scoped to web applications may not display in Central Administration as intended if this service is not started on the server running Central Administration and if the feature cannot be deployed globally.
	Microsoft SharePoint Foundation Workflow Timer Service	Batch-processing server	This service is automatically configured to run on all web servers in a farm.
	Request Management	Distributed Cache and Request Management tier	In integrated mode, Request Management runs on all web servers in a farm. In dedicated mode servers in a separate Request Management farm are between the hardware load balancer and one or more SharePoint farms.

Services on Server for traditional topologies

Table: Service applications and Services on Server for traditional topologies

Service Application	Services on Server	Server Recommendation	More information
Access Services 2010	Access Database Service 2010	Application Server	Allows continued maintenance of SharePoint 2010 Access service applications by using Access 2010 clients and Access 2013 clients. Does not allow users to create new applications.
Access Services	Access Services	Application Server	Allows creation of new Access service applications by using the Access 2013 client. View, edit, and interact with Access Services databases in a browser.
App Management Service	App Management Service	Application Server	
Business Data Connectivity	Business Data Connectivity service	Application Server	
Excel Calculation Services	Excel Calculation Services	Application Server	
Machine Translation Service	Machine Translation Service	Application Server	The Machine Translation Service uses SharePoint OAuth for file management. The Machine Translation Service Application's app pool account needs permissions to the User Profile Application to function correctly. The Machine Translation Service connects to the Microsoft Translator online service. Microsoft Translator may collect some data to improve the quality of translations.
Managed Metadata Service	Managed Metadata Web Service	Application Server	
Microsoft SharePoint Foundation Subscription Settings Service	Microsoft SharePoint Foundation Subscription Settings Service	Web server or application server	This service application is deployed only by using Windows PowerShell. In hosting environments, this service is typically started on one or more application servers. Start this service if you have deployed service applications in multi-tenant mode or if the farm includes sites that use site subscriptions. This service stores settings and configuration data for tenants in a multi-tenant environment. After it is started, web applications consume this service automatically.
PerformancePoint	PerformancePoint	Application server	

	Service		
PowerPoint Conversion	PowerPoint Conversion Service	Application Server	This service converts PowerPoint presentations to other formats. It typically runs on one or more application servers. It starts one or more worker processes to perform conversions. When actively converting, a worker process may use up a whole processor core. Memory usage depends on the size and content of files being converted. The number of worker processes that are used can be controlled through Windows PowerShell cmdlets. Several other configuration options are also available through Windows PowerShell cmdlets.
Search	Lotus Notes Connector	Application Server	Refer to TechNet documentation to learn how to configure Lotus Notes Connector which crawls data on a Lotus Domino server.
Search	Search Host Controller Service	Automatically configured to run on the appropriate application servers	This service manages the search topology components. The service is automatically started on all servers that run search topology components.
Search	Search Query and Site Settings Service	Automatically configured to run on the appropriate application servers	This service load balances queries within the search topology. It also detects farm-level changes to the search service and puts these in the Search Admin database. The service is automatically started on all servers that run the query processing component.
Search	SharePoint Server Search	Automatically configured to run on the appropriate application servers	This service crawls content for the search index. This service is automatically started on all servers that run search topology components. The service cannot be stopped or started from the Services on Server page.
Secure Store Service	Secure Store Service	Application server	
Usage and Health Data Collection	NA	NA	This service application has no associated service on server.
User Profile	User Profile Service	Application server	
User Profile	User Profile Synchronization Service	Application Server	
Visio Graphics Service	Visio Graphics Service	Application server	
Word Automation Service	Word Automation Services	Application server	Performs automated bulk document conversions. When actively converting, this service will fully use one CPU for each worker process (configured in Central Administration). If the service is started on multiple servers, a job will be shared across all the servers.
Work Management	Work Management Service	Application Server	
	Central Administration	Application Server	This service runs the SharePoint Central Administration website.
	Claims to Windows Token Service	Web and application servers	This service is automatically configured to run on applicable servers.
	Distributed Cache	Web and application servers	By default this service is started on all web servers and application servers in a farm.
	Document Conversions Launcher Service	Application Server	Schedules and starts the document conversions on a server.
	Document Conversions Load Balancer Service	Application Server	Balances document conversion requests from across the server farm. Each web application can have only one load balancer registered with it at a time.
	Microsoft SharePoint	Web server or application server	Typically, this service runs on a web server. If you need to isolate this service, you can start it on an application server.

	Foundation Incoming E-Mail		
	Microsoft SharePoint Foundation Sandboxed Code Service	Web server or application server	Start this service on computers in the farm that run sandboxed code. This can include web servers and application servers. This service runs code that is deployed as part of a sandboxed solution in a remote, rights-restricted process and measures the server resources that are used during execution against a site collection-scoped, daily quota.
	Microsoft SharePoint Foundation Web Application	Web server	Ensure that this service is started on all web servers in a farm. Stop this service on application servers. This service provides web server functionality. It is started by default on web servers. Custom features scoped to web applications may not display in Central Administration as intended if this service is not started on the server that runs Central Administration and if the feature cannot be deployed globally.
	Microsoft SharePoint Foundation Workflow Timer Service	Web server	This service is automatically configured to run on all web servers in a farm.
	Request Management	Web server or dedicated servers	In integrated mode, Request Management runs on all web servers in a farm. In dedicated mode servers in a separate Request Management farm are between the hardware load balancer and one or more SharePoint farms.

See also

Plan for SharePoint 2013

SharePoint 2013 design samples: Corporate portal and extranet sites

SharePoint 2013

Applies to: *SharePoint Server 2013, SharePoint Foundation 2013*

Topic Last Modified: *2013-05-08*

Summary: Describes logical architecture design choices for the most common SharePoint 2013 design samples.

This article discusses several design samples that can be used as starting-point architectures for SharePoint sites. The design samples described in this article illustrate standard architectures for the most common types of SharePoint sites that are deployed within a company or organization.

In this article:

Important:
The information in this article applies to SharePoint Foundation 2013 and SharePoint Server 2013. However, the article discusses certain features, such as My Sites and enterprise search, that are not available in SharePoint Foundation 2013.

About the design samples

The following design samples are described in this article:

- Corporate Portal with Path-based Sites
- Corporate Portal with Host-named Sites
- Extranet with Dedicated Zones for Authentication

The design samples illustrate sites for a fictitious company named Fabrikam, Inc. The design samples apply nearly all of the logical architecture components and illustrate how these are incorporated into the overall design. This article describes the design goals for the samples and explains how the logical architecture components illustrated in the samples achieved these goals.

Corporate Portal design sample — two versions

Host-named site collections in SharePoint 2013 provide URL management and scalability of sites within a single web application. The two versions of the Corporate Portal design sample show implementations that are based on the use of the traditional path-based site collections or host-named site collections. Both of these design samples utilize claims-based authentication with a single zone. These samples are discussed in greater detail later in this article.

Extranet with Dedicated Zones for Authentication

The Extranet with Dedicated Zones for Authentication design sample includes only the partner web site. It provides an alternate configuration for partner collaboration in which dedicated zones are used for each authentication method. Each design sample uses claims mode authentication for web applications. The difference between the Corporate Portal design samples and the Extranet Design sample is how the zones are configured. The Extranet with Dedicated Zones for Authentication design sample uses multiple zones, and one method of authentication is configured for each zone. The Corporate Portal design samples use one zone, and multiple authentication methods are configured for different classes of users.

The Extranet with Dedicated Zones for Authentication design sample also introduces a new recommendation for remote employee access — Direct Access with Windows Server 2012. An alternative to this recommendation is to create a virtual private network (VPN). You can also use forms-based authentication on the firewall or gateway product to collect and forward credentials, if desired.

New host-named site collection features

SharePoint 2013 introduces new functionality for host-named site collections (also referred to as host-header site collections). Whereas previous versions of the Corporate Portal design sample made use of path-based site collections, going forward host-named site collections are recommended unless requirements indicate that the traditional path-based sites with alternate access mapping (AAM) are necessary (described later in this article). This is reflected in the design samples in the following ways:

- Corporate Portal with Path-based Site Collections — This sample illustrates path-based site collections in the traditional way with sites organized into dedicated Web applications and a single top-level site collection per Web application. Use this approach if you want the additional security provided by multiple web apps with separate app pools.
- Corporate Portal with Host-named Site Collections — This sample illustrates the use of host-named site collections with all sites deployed in a single Web application on the farm. This method is highly scalable and provides more flexibility in managing URLs.
- Extranet with Dedicated Zones for Authentication — This sample illustrates many top-level project sites with vanity URLs by using host-named sites for each project site (instead of organizing project sites underneath a top-level site collection). One advantage of using host-named site collections in this manner is creating additional isolation between domain URLs, which might be desired in a partner collaboration solution. However, the tradeoff of this approach is the additional costs of managing a greater number of host names, including managing SSL certificates. Also, if SAML authentication is used, additional configuration is required (see "Using SAML authentication with host-named sites" later in this article).

Another design option is retaining a multi-web application design while at the same time implementing host-named site collections. While this option maintains the benefits of using dedicated web applications, it may not be practical because it requires extra configuration work with DNS and IP address binding management within the SharePoint Web applications on each Web server to map requests to the correct Web application when using host-named site collections. However, administrators who are familiar with scripting deployment may choose to manage the configuration of host-named site collections across multiple Web applications by adding the configuration to the scripts. In general, if multiple Web applications are required, the recommendation is to use path-based site collections with AAM instead.

These design options are contrasted later in this article.

Claims-based authentication for SharePoint 2013

The way that authentication works in for SharePoint 2013 might influence design decisions that are related to implementing choices represented by these design samples. Here are some examples:

- In SharePoint 2013, claims-based authentication is the default mode and the only option available through Central Administration. Classic-mode authentication can be implemented by using Windows PowerShell.
- In SharePoint 2013, you don't have to configure server affinity on the load balancer to use SAML claims authentication. SharePoint 2013 fully supports non-affinity load balancing.

In SharePoint 2013 the search crawl account requires access to content through the Default zone by using Integrated Windows authentication (NTLM or Kerberos). Because claims authentication allows multiple types of authentication in one zone, this requirement should not affect other authentication requirements.

Summary of design sample features

The following table summarizes the three design samples that are discussed in this article.

Table: Summary of design samples

Design Sample	Includes	Key design elements
Corporate Portal with Path-based Sites	Most common types of sites deployed within an organization.	• Path-based site collections • Claims-based authentication • Multiple authentication providers and authentication types implemented in a single zone
Corporate Portal with Host-names sites	Most common types of sites deployed within an organization.	• Host-named site collections • Claims-based authentication • Multiple authentication providers and authentication types implemented in a single zone
Extranet with Dedicated Zones for Authentication	Only the partner web site. Provides an alternate configuration for partner collaboration.	• Host-named site collections • Claims-based authentication • Different zone for each authentication method

Sites included in the design samples

This section describes the top-level sites that are included in the design samples.

Intranet sites

The corporate portal includes the following sites for intranet use:

- Published intranet content (such as HRweb)
- Collaborative team sites
- My Sites

Together, these are the content and collaboration sites that employees use on a day-to-day basis. Individually, each of these applications represents a distinct type of content. Each type of content has the following properties:

- Emphasizes different features of SharePoint 2013.
- Hosts data with different data characteristics.
- Is subject to a different usage profile.
- Requires a different permissions management strategy.

Consequently, design choices for each of these applications are intended to optimize the performance and security for each application.

The design of service applications brings these three applications together to provide the following features:

- Enterprise-wide search
- Shared profile data and enterprise metadata

The following illustration, from the Corporate Portal with Path-based Site Collections design sample, shows the three types of sites that make up the corporate intranet.

Types of sites that make up the corporate intranet

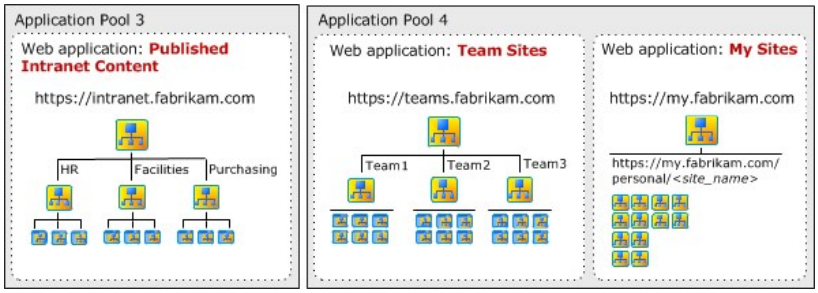

Partner web application

The partner web application hosts externally-available sites for secure collaboration with partner companies and individual partners. This application is intended for employees to easily create sites for secure collaboration. Partners cannot access other types of content that the server farm hosts. The design for zones and service applications addresses this goal. Additionally, individual site owners manage permissions for their sites and invite only necessary participants to collaborate.

In the extranet design sample, this is the only type of represented site.

Overall design goals

The design samples provide practical implementations of SharePoint 2013 features within several common types of sites. The design implementations for each of the individual applications are discussed in this article. Following are the key design goals for the design samples:

- Create a framework for designing an environment that can grow.

  Design decisions for individual types of sites do not prevent the addition of other types of sites. For example, an initial deployment might include only collaborative team sites or only the three types of sites that compose an intranet (team sites, My Sites, and published intranet content). If you use a similar logical architecture design, you can add sites to the solution without affecting the design of the initial solution. In other words, the design does not incorporate design choices that limit the use of the environment.

- Provide access for several groups of users without compromising the security of the content within the different types of sites.

Users from different network zones (both internal and external) who use different authentication providers can participate in collaboration. Also, users can only access the content they are intended to access. If you follow a similar logical architecture design, you can provide access to users who are in multiple locations and have different objectives. For example, your initial design might be intended only for internal employee access. However, if you use a similar design, you can also enable access to remote employees, partner employees, partner companies, and customers.

- Ensure that the design can be used in an extranet environment.

  Make deliberate design choices to make sure that the solution can be securely deployed in a perimeter network.

The rest of this article discusses each of the logical components that appear in the design sample (from top to bottom) and discusses the design choices that are applied to the design sample. The purpose of this approach is to demonstrate the different ways in which logical architecture components can be configured based on the application.

Server farms

This section describes the topologies of the server farms that are illustrated in the design sample and discusses scaling beyond a single farm.

Topology of the server farms

Each server farm in the design sample is composed of six servers with the following fault-tolerant topology:

- Two front-end web servers
- Two application servers
- Two database servers with SQL Server installed and configured to support SQL Server clustering, mirroring, or AlwaysOn. AlwaysOn requires SQL Server 2012.

The design sample illustrates the logical architecture of SharePoint 2013 by showing that the following:

- All sites are mirrored across front-end web servers.
- The Central Administration site is installed on an application server to protect it from direct user access.

In reality, the number of server computers and the topology of the server farm are important to the logical architecture only to increase capacity and improve performance. You can design the logical architecture independent of the server farm topology. The process of planning performance and capacity helps you plan the size the server farm to meet performance and capacity goals.

Scale beyond two farms

Your business might require more than one farm. Sites that are candidates for a dedicated farm include the following:

- My Sites

  Many organizations with large numbers of employees or students choose to host My Sites on a dedicated server farm.

- Partner sites

  Security and isolation requirements might warrant a dedicated farm for partner collaboration. This physically isolates internal-only content from content that is developed in collaboration with external partners. The extranet design sample illustrates this configuration.

Users, zones, and authentication

Claims is the default mode of authentication in SharePoint 2013 and each design sample incorporates claims-based authentication. You can use Windows PowerShellto implement classic-mode authentication; however, some features of SharePoint 2013 are not available in classic mode. For more information about design samples that incorporate classic-mode authentication, see Design sample: Corporate deployment (SharePoint Server 2010)

The following table lists the differences in the two approaches that are represented by the Corporate Portal Design sample and the Extranet Design sample.

Table: Contrasting approaches for zone configurations in the design samples

	Corporate Portal with Host-named sites and Corporate Portal with Path-based sites	**Extranet with Dedicated Zones for Authentication**
Mode of authentication	Claims	Claims
Zone configuration	One zone with multiple authentication methods configured for different classes of users.	Multiple zones with one method of authentication configured for each zone.
URLs	All classes of users use the same URL for each site. Employees use the same URL regardless of whether they are inside the corporate network or working remotely.	Each class of user uses a different URL to access a site. Employees use a different URL depending on whether they are

		inside the corporate network or working remotely.
Internal requests	Requests that initiate inside the corporate network are routed out of the network and then back in through the gateway or proxy server. These requests are secured by using the Secure Sockets Layer (SSL) protocol. Alternatively, split DNS can be used to route the requests directly to the internal interface for the servers.	Requests that initiate inside the corporate network remain internal to the network.
User experience	All users are prompted to choose the type of account they are using to log in.	The authentication method is predetermined. Users are not required to select the account type by using a logon page.

The following sections specifically discuss how authentication is incorporated by using the two different approaches.

Extranet with dedicated zones design sample

The extranet design sample illustrates three different classes of users, and each is assigned to a different zone. Within each web application, you can create up to five zones by using one of the available zone names: Default, Intranet, Internet, Custom, or Extranet. The search crawl account requires access to the Default zone by using Integrated Windows authentication (NTLM or the Kerberos protocol), which is accounted for in the design sample. The following table shows how zones are set up in the extranet design sample.

Table: Zones, users, and authentication type prescribed by the extranet design sample

Zone	Users	Authentication
Intranet	Internal and remote employees Search crawl account	NTLM or Kerberos protocol Remote employees who use Direct Access or VPN to connect.
Default	Individual partners	Options: • LDAP directory using forms-based authentication • Externally-facing Active Directory Domain Services (AD DS) forest with a one-way trust to the internal forest and Integrated Windows authentication • Trusted identity provider with SAML authentication
Extranet	Partner companies	Trusted partner identity provider with SAML authentication

Corporate Portal design samples

Claims-based authentication allows multiple types of authentication in the same zone. The two versions of the Corporate Portal design sample use the Default zone for all authentication types.

The following table shows the zones, users, and authentication types that are prescribed by the design samples.

Table: Zones, users, and authentication for the corporate portal design samples

Zone	Users	Provider and authentication type
Default	Internal and remote employees	Active Directory Domain Services (AD DS) or LDAP store with forms-based authentication or SAML authentication.
Default	Individual partners	Trusted identity provider with SAML authentication, or SQL Server database with forms-based authentication
Default	Partner companies	Trusted partner identity provider with SAML authentication
Default	Search crawl account	AD DS with Windows NTLM authentication

In the design sample, the Published Intranet Content site, Team Sites, and My Sites are only accessible to employees, whether they are inside or outside the network. The design sample implements only one URL (using SSL) for each of these sites that can be used both internally and externally. AD DS accounts are used. If needed, forms-based authentication or SAML can use LDAP, which requires additional configuration.

In the design sample, the partner web application represents an extranet site that partner employees and partner companies can access. Claims-based authentication in this scenario requires you to configure trust with one or more external identity providers. You can use either one of the following approaches:

- You can configure the SharePoint farm to trust an external identity provider, such as the provider that resides in a partner company (to authenticate directly against the partner directory).
- You can configure the identity provider inside the corporate environment to trust an external identity provider. Administrators in the two organizations must establish this relationship explicitly. In this scenario, the SharePoint farm trusts the identity provider from within its own corporate environment. When the identity provider sends a token, the farm uses the token signing certificate that was specified when the trust was established to confirm the validity of the token. We recommend this approach.

Forms-based authentication is an alternative to a claims-based environment to authenticate partners. You use a separate store, such as a database, to manage these accounts.

Zones

When you design zones, several key decisions are critical to the success of the deployment. These decisions include design and configuration decisions for the following zones:

- The Default zone
- Zones for external access

The following sections describe the decisions that are incorporated in the design sample.

Configuration requirements of the default zone

The zone that involves the greatest consideration is the Default zone. SharePoint 2013 places the following requirements on how you configure the Default zone:

- When a user request cannot be associated with a zone, the authentication and policies of the Default zone are applied. Consequently, the Default zone must be the most secure zone.
- Administrative e-mail messages include links from the Default zone. These include messages to owners of sites that are approaching quota limits. Consequently, users who receive these kinds of messages and alerts must be able to access links through the Default zone. This is especially important for site owners.

In SharePoint 2013 host-named site collections can be access from any zone.

Configuring zones for an extranet environment

In an extranet environment, the design of zones is critical for the following two reasons:

- Several different networks can initiate user requests. In the design samples, users initiate requests from the internal network, the Internet, and partner companies.
- Users consume content across multiple web applications. In the design sample, the intranet is composed of three web applications. Additionally, internal and remote employees can potentially contribute to and administer content in the partner web application.

If an extranet environment includes more than one zone, ensure that you follow these design principles:

- Configure zones across multiple web applications to mirror each other. The configuration of authentication and the intended users should be the same. However, policies that are associated with zones can differ across web applications. For example, make sure that you use the Intranet zone for the same employees across all web applications. In other words, do not configure the Intranet zone for internal employees in one web application and remote employees in another.
- If you use path-based site collections, configure alternate access mappings appropriately and accurately for each zone and each resource. Alternate access mappings are automatically created when you create a zone. However, you can configure SharePoint 2013 to crawl content in external resources, such as a file share. You must create links to these external resources manually for each zone by using alternate access mappings.
- If you use host-named site collections, make sure that you use Windows PowerShell to map URLs to the appropriate zones

If zones across Web applications do not mirror each other and links to external resources are not appropriate, the following risks can occur:

- Server names, Domain Name System (DNS) names, and IP addresses can potentially be exposed outside the internal network.
- Users might be unable to access Web sites and other resources.

Using SAML authentication with host-named sites

If a design includes the use of SAML authentication with host-named sites, each vanity URL requires the following:

- A new realm on the **SPTrustedIdentityTokenIssuer**
- A corresponding relying party in the Identity Provider.

Services

The service architectures vary depending on the design sample. The Corporate Portal with Host-named Sites design sample includes the simplest architecture for

services. This is because it uses one web application, which can accommodate only one service application group (also known as proxy group).

The Corporate Portal with Path-based Sites example uses partitioned services for the Partner sites to isolate data between projects. This design sample incorporates two service groups, one for the intranet sites and one for the partner collaboration sites. Separate instances of Managed Metadata and Search are deployed for the Partner sites and these services are partitioned. Partitioned services require the Subscription Settings Service which can only be deployed by using Windows PowerShell.

Service architecture for the Corporate Portal with Path-based Sites

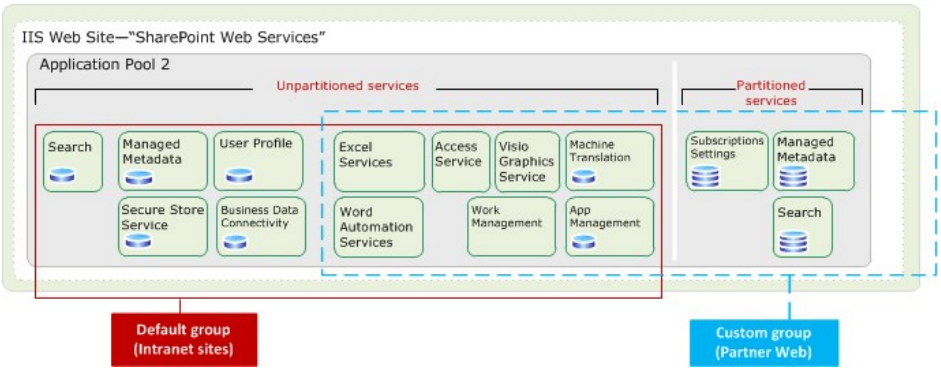

Deploying partitioned services adds complexity to the architecture and makes it difficult to migrate sites to Office 365 at a later time. A simpler option for the Partner sites is to deploy dedicated but unpartitioned instances of the Managed Metadata service and Search service if these are required to be separate. Many organizations rely on the security trimming feature of Search, rather than deploying dedicated instances of the Search service.

The Extranet design sample includes only one proxy group but it also uses partitioned services for both the Managed Metadata and Search service applications.

The primary design decision for deploying service applications is how broadly to spread the organization taxonomy. To simplify the services architecture, share Managed Metadata, User Profile, and Search across all web apps and rely on security trimming to manage access to content. In the Corporate Portal with Path-based Sites design sample, one instance of the Managed Metadata service is shared across all sites. However, all users can access to the corporate taxonomy with this configuration. Solution architects must decide whether to implement multiple instances of the Managed Metadata service. They'll also need to decide how broadly to share the User Profile data.

Within the Corporate Portal with Path-based Site Collections sample, partner web is configured to use dedicated Search and Managed Metadata service applications by using a custom service group. Other service applications are added to the custom group and these are shared with the default group. The design sample does not include the User Profile service application to prevent partner users from browsing people data in the organization.

In the simplified architecture of the Corporate Portal with Host-named Site Collections (one service group), partners can access the entire corporate taxonomy and can browse people data in the organization. However, search limits results to sites and content that partners can access.

If your partner sites require content isolation between projects, deploying dedicated service applications is a good choice, as illustrated in this article. This increases the complexity of the services architecture but ensures that partners cannot access metadata that is associated with the Intranet content or even other projects within the partner web site. The Extranet design sample also uses partitioned services.

Administration sites

In the design sample, an application server hosts the SharePoint Central Administration website for each server farm. This protects the site from direct user contact. If a performance bottleneck or security compromise affects the availability of the front-end web servers, the SharePoint Central Administration website remains available.

The design sample and this article do not mention the load-balanced URLs for administration sites. Recommendations include the following:

- If administrative URLs use port numbers, use non-standard ports. URLs include port numbers by default. While customer-facing URLs typically do not include port numbers, using port numbers for administration sites can increase security by limiting access to these sites to non-standard ports.
- Create separate DNS entries for administration sites.

In addition to these recommendations, you can optionally load-balance the SharePoint Central Administration website across multiple application servers to achieve redundancy.

Application pools

Separate Internet Information Services (IIS) application pools are typically implemented to achieve process isolation between content. Application pools provide a way for multiple sites to run on the same server computer but still have their own worker processes and identity. This helps to prevent an attacker who injects code on one site from affecting other servers or sites on other sites.

If a single application pool and web application is used together with host-named site collections, isolation is provided between domain URLs but only at the scripting level.

If you choose to implement more than one application pool, consider a dedicated application pool for each of the following scenarios:

- To separate authenticated content from anonymous content. If the same farm hosts the company Internet site, place this site in a dedicated web application and application pool.
- To isolate sites that store passwords for and interact with backend data systems (although the Secure Store Service can be used for this purpose instead).

The Corporate Portal with host-named sites design sample and the Extranet with Dedicated Zones for Authentication design sample both implement a single application

pool and web application for all content. Separate application pools are required for service applications and the SharePoint Central Administration website.

The Corporate Portal with Path-based Sites design sample implements process isolation between content by using separate application pools in the following ways:

- The administration site is hosted in a dedicated application pool. This is a requirement of SharePoint 2013.
- All service applications are deployed to a single application pool. Unless there is a compelling reason to deploy service applications to different application pools, this is the recommended configuration. One application pool for all service applications optimizes performance and reduces the number of application pools to manage.
- Intranet content is divided into two different application pools. One application pool hosts collaborative content (My Sites and team sites). A separate application pool hosts the published intranet content. This configuration provides process isolation for the published intranet content in which business data connections are more likely to be used.
- A dedicated application pool hosts the partner web application.

Web applications

A web application is an IIS web site that SharePoint 2013 creates and uses. Each Web application is represented by a different Web site in IIS.

If you choose to implement more than one web application, consider the following use cases:

- **Separate anonymous content from authenticated content.**

  If the same farm hosts the company Internet site, place this site in a dedicated web application and application pool.

- **Isolate users**

  In the design sample, a dedicated web application and application pool hosts the partner web site to make sure that partners do not have access to the intranet content.

- **Enforce permissions**

  A dedicated web application provides the opportunity to implement policies on the zones within the web application to enforce permissions. For example, you can create policies for the company Internet site to explicitly deny write access to one or more groups of users. Policies are enforced regardless of permissions that are configured on individual sites or documents in the web application.

- **Optimize performance**

  Applications achieve better performance if you place them in web applications with other applications of similar data characteristics. For example, the data characteristics of My Sites include a large number of sites that are small in size. In contrast, team sites typically encompass a smaller number of very large sites. By placing these two different types of sites in separate web applications, the resulting databases are composed of data with similar characteristics, which optimizes database performance. In the design sample, My Sites and team sites do not have unique data isolation requirements—they share the same application pool. Nonetheless, My Sites and team sites are placed in separate web applications to optimize performance.

- **Optimize manageability**

  Because creating separate Web applications results in separate sites and databases, you can implement different site limits (recycle bin, expiration, and size) and negotiate different service-level agreements. For example, you might allow more time to restore My Site content if this is not the most critical type of content within your organization. This enables you to restore more critical content before you restore My Site content. In the design sample, My Sites are in a separate web application to enable administrators to more aggressively manage growth compared to other applications.

If you implement host-named site collections with a single web application, each top-level site is a separate domain which enables you to achieve some of these goals, such as optimizing manageability by implementing different site limits.

Site collections

The recommended configuration for deploying sites is using host-named site collections with all sites located within a single web application. This configuration is recommended to deploy sites because it is the same architecture that the Office 365 environment uses. Consequently this is the most heavily tested configuration. New features, including the App model and Request Management, are optimized for this configuration, and it is the most reliable configuration going forward.

Although we recommend host-named site collections for most architectures, you should use the traditional path-based site collections and alternate access mapping if any of the following conditions apply:

- You need to use the Self Service Site Creation feature that is part of the default installation of SharePoint 2013.

  This does not apply to custom self-service site creation solutions.

- A web application requires unique wild card inclusions or explicit inclusions.

  You create inclusions for host-named site collections at the farm level, and they are available for all host-named sites. All host-named site collections in a farm share the same inclusions but do not need to use them. In contrast, inclusions that you create for path-based site collections apply only to a single web application.

- SSL termination is required but your SSL termination device cannot be configured to produce the necessary custom HTTP header.

  You can still use SSL bridging with host-named site collections with these devices if SSL termination is not a requirement.

- You plan to use different application pools for the additional security that these provide or you need to use multiple proxy groups.

For more information about host-named site collections, including a comparison with path-based site collections, see Host-named site collection architecture and deployment (SharePoint 2013).

Design goals for site collections

Site collections bridge logical architecture and information architecture. The design goals for site collections are to fulfill requirements for URL design and to create logical divisions of content. For each site collection, managed paths incorporate a second tier of top-level site collections. For more information about URL requirements and using managed paths, see Zones and URLs later in this article. Beyond the second tier of site collections, each site is a subsite.

The following diagram illustrates the site hierarchy of Team Sites.

Site hierarchy for team sites

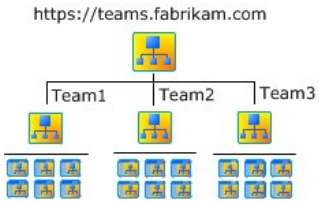

For both path-based site collections and host-named collections, the information architecture revolves around the second tier of site collections. The following section describes how design samples incorporate choices based on the nature of the sites.

Published intranet content

The assumption for the published intranet content web application is that multiple divisions within the company will host published content. In the design sample, a separate site collection hosts each division's content. This provides the following advantages:

- Each division can manage content and administer permissions independently.
- Each division can store its content in a dedicated database.

The disadvantages of multiple site collections include the following:

- You cannot share master pages, page layouts, templates, Web Parts, and navigation be across site collections.
- Coordination of customizations and navigation across site collections requires more effort.

Depending on the information architecture and design of the intranet sites together, the published content can appear to the user as a seamless experience. Alternatively, each site collection can appear to be a separate web site.

My Sites

My Sites have distinct characteristics and the deployment recommendations for My Sites are straightforward. In the design samples, the My Sites site collection incorporates a top-level site with the URL of http://my. The first top-level site collection that is created uses the My Site Host template. A managed path is incorporated (by using wildcard inclusion), which allows an indefinite number of user-created sites. All sites below the managed path are independent site collections that use the Personal Site template. The user name is appended to the URL in the form http://my personal/username. The following illustration illustrates My Sites.

Site hierarchy for My Sites

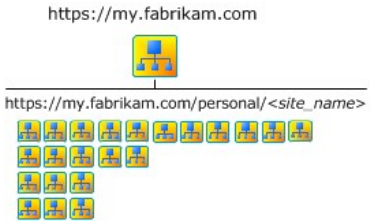

Team sites

You can use either of the following two approaches to design site collections in a Team Site application:

- Allow teams to create site collections through self-service site creation. The advantage of this approach is that teams can easily create a site, as needed, without assistance from an administrator. Disadvantages to this approach include the following:
  - You lose the opportunity to implement a thoughtful taxonomy.
  - You cannot share templates and navigation across projects or teams that might otherwise share a site collection.
- Create a finite number of site collections for your organization based on the way your organization operates. In this approach, a SharePoint administrator creates site collections. After a site collection is created, teams can create sites within the site collection. This approach provides the opportunity to implement a thoughtful taxonomy that provides structure to the way team sites are managed and grow. There is also more opportunity to share templates and navigation between projects and teams that share a site collection. However, this approach also includes some disadvantages.

The design samples incorporate the second approach, which results in a similar site collection hierarchy for team sites and published intranet content. The challenge for information architects is to create a second tier of site collections that makes sense for the organization. The following table suggests different types of organizations.

Table: Suggested site collection taxonomies

Type of organization	Suggested site collection taxonomies
Product development	• Create a site collection for each product under development. Allow contributing teams to create sites within the site collection. • For each long-term development project, create a site collection for each large team that contributes to the product. For example, create one site collection for each of the following teams: designers, engineers, and content developers.
Research	• Create a site collection for each long-term research project. • Create a site collection for each category of research projects.
Higher education institution	• Create a site collection for each academic department.
State legislative office	• Create a site collection for each political party. Government officials who share party affiliation can share templates and navigation. • Create a site collection for each committee. Or, create one site collection for all committees.
Corporate law office	• Create a site collection for each corporate client.
Manufacturing	• Create a site collection for each line of products.

Partner web application

Partner web is intended to be used for collaboration with external partners on projects that have finite scopes or finite durations. By design, sites in the partner web application are not intended to be related. The requirements for the partner web application include ensuring that:

- Project owners can easily create sites for partner collaboration.
- Partners and other contributors can access only the project they work on.
- Site owners manage permissions.
- Search results from within one project do not expose content from other projects.
- Administrators can easily identify sites that are no longer used and delete these sites.

To satisfy these requirements, the design sample incorporates a site collection for each project. Both Corporate Portal design samples utilize a managed path to create a second tier of site collections below a root site collection. Alternatively, the Extranet design sample makes each partner site a top-level site collection by using host-named site collections. Either way, individual site collections provide the appropriate level of isolation between projects.

If you plan to use the Self Service Site Creation feature that is part of the default installation of SharePoint 2013 (as opposed to a custom solution developed for your organization), then use path-based site collections. Host-named site collections do not yet work with this feature.

Content databases

You can use the following two approaches to incorporate content databases into the design (the design sample incorporates both approaches):

- Establish target sizes for content databases with appropriate size-warning thresholds. Create a new database when a database reaches size-warning thresholds. With this approach, site collections are automatically added to the available database or databases, based on size targets alone. This is the most commonly used approach.
- Associate site collections to specific content databases. This approach enables you to place one or more site collections in a dedicated database that administrators can manage independently from the rest.

If you choose to associate site collections to specific content databases, you can use the following methods to accomplish this:

- Use Windows PowerShell to create a site collection in a specific database.
- Dedicate a database to a single site collection by applying the following database capacity settings:
    - Number of sites before a warning event is generated = 1

    - Maximum number of sites that can be created in this database = 1

- Add a group of site collections to a dedicated database by completing the following steps:

    1. Within the web application, create the database and set the database status to **Ready**.
    2. Set the status of all other databases to **Offline**. While content databases are offline, new site collections cannot be created. However, existing site collections in offline databases are still accessible for both read and write operations.
    3. Create the site collections. They are automatically added to the database.
    4. Set the status of all other databases back to **Ready**.

Published intranet content

For published intranet content, the corporate portal design samples incorporate a single database for ease of management. Add databases based on target size goals, if needed.

My Sites

For My Sites, the corporate portal design samples achieve scale efficiency by managing databases to the maximum target size. The following settings are configured to achieve this goal:

- **Limit site storage to a maximum of**: This setting, which you configure on the **Quota Templates** page in Central Administration, limits the size of a personal site.
- **Second stage Recycle Bin**: This setting, which you configure on the **Web Application General Settings** page, determines the amount of additional space that is allocated to the second-stage recycle bin.
- **Maximum number of sites that can be created in this database**: This setting is configured when you create a database. Calculate the total allowable size of sites by using the numbers you specify for the previous two values. Then, based on the size goal for each database, determine how many sites will fit in the database.

The design samples provide the following example size settings based on a target database size of 175 gigabytes (GB) and a target My Site size of 1 GB:

- Site size limits per site = 1 GB
- Target size of database = 175 GB
- Reserved for second-stage recycle bin = 15%
- Maximum number of sites = 180
- Site level warning = 150

When the site-level warning is reached, create a new database. After you create the new database, new My Sites are added to the content database that has the fewest site collections.

Team sites

For team sites, again, the corporate portal design samples achieve scale efficiency by managing databases to the maximum target size. Team sites for most organizations are expected to be much larger than My Sites. The design sample provides database settings based on a 30-GB limit for site collections. Choose a limit that is appropriate for team sites in your organization.

Another approach for organizations whose teams have large storage needs is to dedicate a single database to each top-level team site collection.

Partner web

Similar to My Sites, partner web achieves scale efficiency by managing databases to the maximum target size.

The design samples provide the following example size settings:

- Target size of database = 200 GB
- Storage quota per site = 5 GB
- Maximum number of sites = 40
- Authoring site collection hosted in dedicated database

Zones and URLs

The design samples illustrate how to coordinate URLs across multiple sites within a corporate deployment. The following goals influence design decisions for URLs:

- URL conventions do not limit the zones through users can access content.
- The standard HTTP and HTTPS ports (80 and 443) can be used across all applications in the design sample.
- Port numbers are not included in URLs. In practice, port numbers are typically not used in production environments.

Designing load-balanced URLs

When you create a web application, you must choose a load-balanced URL to assign to the application. The URL that you choose applies to the Default zone. Additionally, you must create a load-balanced URL for each additional zone that you create within a web application. The load-balanced URL includes the protocol, scheme, hostname, and port, if used. The load-balanced URL must be unique across all web applications and zones. Consequently, each web application and each zone within each web application requires a unique URL across the design sample.

Intranet

Each of the three site collections that make up the intranet requires a unique URL. In Corporate Portal design samples, the target audience for the intranet content is internal employees and remote employees. Employees use the same URLs for each of these sites regardless of whether they are on site or remote. This approach adds a layer of security to the SharePoint design (all traffic is SSL). However, this approach requires you to choose an alternative for additional configuration:

- Route internal traffic through the firewall or gateway product along with remote traffic.
- Set up a split DNS environment to resolve internal requests within the internal network.

Partner web site

In the design samples, internal employees, remote employees, and partner employees access the partner web site. In the Corporate Portal design samples, all users enter the same URL regardless of the authentication method. In the Extranet design sample, each different type of user enters a different URL. Although both individual partners and partner companies use SSL (HTTPS) to access the partner web site externally, each group requires a different URL to apply the benefits of separate zones—that is, different authentication methods and different zone policies.

Because the Extranet design sample uses Direct Access or VPN for remote employee access, both remote employees and internal employees use the same URLs. If access for remote employees is configured through a reverse proxy device, remote employees would require a separate URL using SSL, requiring an additional zone. Finally, the Extranet design sample incorporates host-named site collections instead of a single top-level site collection. Consequently, each project site has a different URL.

The following table shows example URLs that internal employees, remote employees, and partners use to access the partner web site, as shown in the Extranet design sample.

Table: Example URLs from the Extranet design sample

Zone	Example URL
Internal and remote employees	http://project1
Individual partners	https://project2.fabrikam.com
Partner companies	https://TrustedPartnerProject1.fabrikam.com

Using explicit and wildcard inclusions for URL paths

Managed paths enable you to specify the paths in the URL namespace of a web application that are used for site collections. You can specify that one site collection or more than one site collection exists at a distinct path below the root site. Without managed paths, all sites below the root site collection are part of the root site collection.

You can create the following two types of managed paths:

- **Explicit inclusion:** A site collection with the explicit URL that you assign. An explicit inclusion is applied to only one site collection. You can create many explicit inclusions below a root site collection. . An example URL for a site collection created by using this method is http://intranet/hr. There is a performance impact for every explicit path added so the recommendation is to limit the number of site collections created with an explicit inclusion to about 20.
- **Wildcard inclusion:** A path that is added to the URL. This path indicates that all sites that are specified directly after the path name are unique site collections. This option is typically used for site collections that support self-site creation, such as My Sites. An example URL for a site collection that is created by using this method is http://my/personal/user1.

When managed paths for host-named site collections are implemented, these managed paths are created at the farm level and paths apply across all web applications, if multiple web applications are included in the solution. When managed paths for path-based sites are implemented, these managed paths apply only to the Web application in which they were created.

The design sample incorporates the use of both types of managed paths (explicit inclusions and wildcard inclusions, as described in the following sections.

Explicit inclusions: Published intranet content

In the design samples, the published intranet site collection incorporates explicit inclusions for each subsite, for example, HR, Facilities, and Purchasing. Each of these site collections can be associated with a different content database, if needed. Unless host-named site collections are used, the use of explicit inclusions in this example assumes that no other types of sites are created in the web application, including wildcard inclusions.

The use of explicit inclusions results in the URLs:

- https://intranet.fabrikam.com
- https://intranet.fabrikam.com/hr
- https://intranet.fabrikam.com/facilities
- https://intranet.fabrikam.com/purchasing

In this example, the root site collection, http://intranet.fabrikam.com, represents the default home page for the intranet. This site is intended to host content for users.

Wildcard inclusions: Team Sites, My Sites, and Partner Web

Team Sites, My Sites, and the partner web application incorporate the use of a wildcard inclusion. Wildcard inclusions are ideal for applications that allow users to create their own site collections and for web applications that include many site collections. A wildcard inclusion indicates that the next item after the wildcard is a root site of a site collection.

Team sites

Within the Team Sites application, wildcard inclusion is used for each team site collection. Good governance practices recommend that you keep the number of top-level team sites within a manageable number. Also, the taxonomy for team sites should be logical for the way your business operates.

The use of wildcard inclusions results in the URLs:

- https://teams.fabrikam.com/sites/Team1
- https://teams.fabrikam.com/sites/Team2
- https://teams.fabrikam.com/sites/Team3

In this example, the root site collection, https://teams.fabrikam.com, does not necessarily host content for users.

My Sites

My Sites offer self-service site creation. When a user who browses the intranet first clicks **My Site**, a My Site is automatically created for the user. In the design sample, My Sites include a wildcard inclusion named /personal (http://my/personal). The My Site feature automatically appends the user name to the URL.

This results in URLs of the format:

- https://my.fabrikam.com/personal/User1
- https://my.fabrikam.com/personal/User2
- https://my.fabrikam.com/personal/User3

Partner web

If path-based site collections are used, you can implement the self-service site creation feature to allow employees to create secure sites for collaboration with external partners. If host-named site collections are used, you can implement a custom self-service site creation feature or administrators can create partner project sites by request.

In Corporate Portal design samples, the partner web application includes a wildcard inclusion named /sites (http://partnerweb/sites). This results in URLs of the following format:

- https://partnerweb.fabrikam.com/sites/Project1
- https://partnerweb.fabrikam.com/sites/Project2
- https://partnerweb.fabrikam.com/sites/Project3

Coordinating URLs with AAM and DNS

If path-based site collections are implemented, configure alternate access mappings (AAM) for each site URL in the farm. This makes sure that web requests are mapped to the correct site, especially in environments that use load balancing or reverse proxy technologies.

Single-name URLs, such as http://teams, can be configured for intranet access. A client computer resolves these URLs by appending the DNS suffix of the client computer, such as fabrikam.com, and then issuing a DNS lookup for the name with the suffix. For example, when a client computer in the fabrikam.com domain requests http://teams, the computer sends a request to DNS for http://teams.fabrikam.com.

DNS must be configured to use an A record, or AAAA for IPv6, for each fully qualified domain name (FQDN). The record points to the load-balanced IP address for the web servers that host a site. In a typical production deployment, servers are configured to use statically assigned IP addresses, in addition to statically assigned A or AAAA records in DNS.

After a client browser receives the load-balanced IP address, the client browser connects to a front-end web server in the farm, and then sends an HTTP request that has the original single-name URL, http://teams. IIS and SharePoint 2013 recognize this as a request for the Intranet zone, based on the settings that are configured in alternate access mappings. If a user instead requests https://teams.fabrikam.com, the process is similar, but IIS and SharePoint 2013 receive this FQDN instead, and therefore recognize this request for the Default zone.

In environments that have multiple domains, enter CNAME records for DNS in the domains that the sites do not reside in. For example, if the Fabrikam network environment includes a second domain named europe.fabrikam.com, CNAME records are entered for these sites in the Europe domain. For the Team Sites intranet site (http://teams), a CNAME record named teams is added to the europe.fabrikam.com domain that points to teams.fabrikam.com. Then, when a client computer's DNS suffix is appended to DNS lookup requests, a request for http://teams from the Europe domain will issue a DNS lookup of teams.europe.fabrikam.com, and will be directed by the CNAME record to teams.fabrikam.com.

Note:

There is a known issue with some clients that use Kerberos authentication and resolving CNAME records. For more information, see Kerberos configuration known issues (SharePoint Server 2010).

Zone policies

You can configure policies for one or more zones to enforce permissions for all content within a web application. In claims mode, a policy can be defined only for a specific zone (not for the web application in general). A policy enforces permissions on all content that users access through a zone. Policy permissions override all other security settings that are configured for sites and content. You can configure policy based on users or user groups, but not SharePoint groups. If you add or change a zone policy, search must crawl sites again to apply the new permissions.

The design samples do not use policies because either multiple types of authentication are enabled on a single zone or all sites are contained within on web application (or both).

Host-named site collection architecture and deployment (SharePoint 2013)

SharePoint 2013

Applies to: *SharePoint Server 2013, SharePoint Foundation 2013*

Topic Last Modified: *2013-12-18*

Summary: Plan and implement host-named site collections in SharePoint 2013 and learn how path-based site collections might affect your environment.

Host-named site collections are the preferred method to deploy sites in SharePoint 2013. Because the Office 365 environment uses host-named site collections, new features are optimized for these site collections and they are expected to be more reliable. Learn how to plan for and implement host-named site collections, design URLs, and manage URLs.

In this article:

Architecture and design for host-named site collections

Host-named site collections enable you to assign a unique DNS name to site collections. For example, you can address them as http://TeamA.contoso.com and http://TeamB.contoso.com. This enables you to deploy many sites with unique DNS names in the same web application. It also enables hosters to scale an environment to many customers. If you do not use host-named site collections, your SharePoint web application will contain many path-based site collections that share the same host name (DNS name). For example, Team A has a site collection at http://contoso.com/sites/teamA, and Team B has a site collection at http://contoso.com/sites/teamB.

We recommend host-named site collections unless requirements dictate that path-based sites with alternate access mapping are necessary (described later in this article). This article describes how to implement host-named site collections in a recommended configuration with SharePoint 2013. Information about advanced configurations is included at the end of this article: Use multiple web applications with host-named site collections.

Recommended architecture for host-named site collections

The recommended configuration for deploying sites is using host-named site collections with all sites located within a single web application, as illustrated in the following diagram.

Recommended configuration for host-named site collections

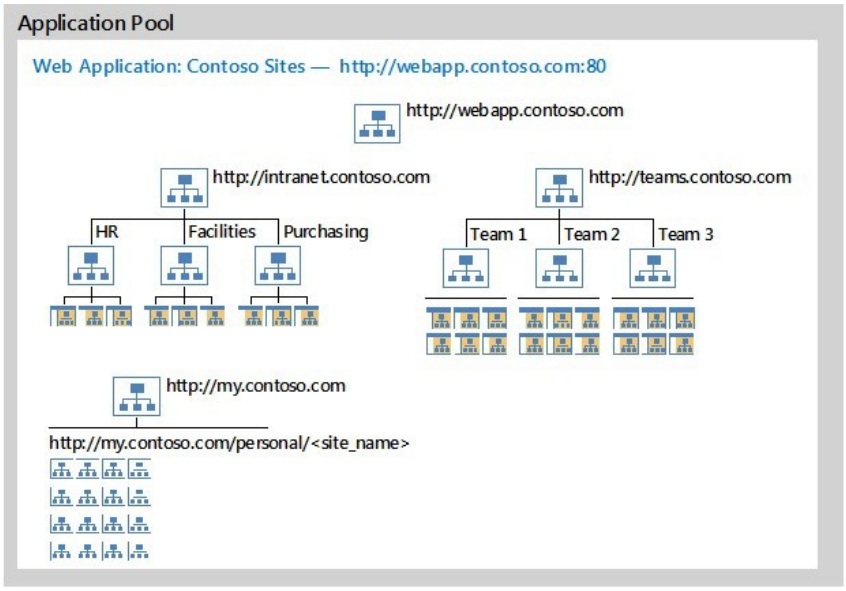

This recommended configuration in the diagram includes the following elements:

- One application pool for site collections.
- One web application for site collections that is hosted inside the one application pool.
- A root site collection (http://SP1).
- Multiple host-named site collections to host content with example sites:
  - Published intranet content (http://intranet.contoso.com) with subsites for HR, Facilities, and Purchasing.
  - Team sites (http://teams.contoso.com) with subsites for Team 1, Team 2, and Team 3.
  - My Sites with site URLs in the following format: http://my.consoso.com/personal/<site_name>.

The number of sites within the web application and the URLs for sites are not important for this example.

When creating a Web application for host-named site collections, the URL of the Web application and the root site collection will be http://< *servername*>/.

This architecture is recommended to deploy sites because it is the same architecture that the Office 365 environment uses. Consequently this is the most heavily tested configuration. New features, including the App model and Request Management, are optimized for this configuration, and it is the most reliable configuration going forward.

The recommended configuration does not include the following elements:

- Enabling apps in environments with multiple zones.
- Mixing host-named site collections and path-based site collections (except the root site collection).
- Multiple web applications with host-named site collections.

Host-named site collections versus path-based site collections

When you use host-named site collections, each site collection in a web application is assigned a unique DNS name. When you deploy many host-named site collections to a single web application, you increase scalability of the farm because resources are not used to support multiple application pools and web applications.

SharePoint 2013 supports both host-named and path-based site collections. The following table details the differences between the two options and provides more information about host-named site collections.

Table: Comparison of host-named site collections and path-based site collections

	Host-named site collections	Path-based site collections
Creating sites	You can use Windows PowerShell to create host-named site collections. You cannot use Central Administration to create host-named	You can use Central Administration or Windows PowerShell to create path-based site collections.

	site collections.	
URLs	Each host-named site collection in a web application is assigned a unique DNS name. You can use zones to assign up to five URLs to host-named sites, including vanity URLs.	All path-based site collections in a web application share the same host name (DNS name) as the web application. You can extend a web application to implement up to five zones and create different host names for each zone. However, the host name for a zone applies to all site collections within the web application.
Root site collection and search	A root site collection is required to crawl content in a web application. A root site collection can be a site collection that users cannot access.	Typically, a single path-based site collection serves as the root site collection within a web application. You can use managed paths to create additional site collections within the web application.
URL mapping	Use Windows PowerShell commands to manage URLs (Set-SPSiteURL, Remove-SPSiteURL, Get-SPSiteURL).	Use Alternate Access Mappings to manage URLs.
Self-service site creation	You need to use a custom solution for self-service site creation with host-named site collections. The Self Service Site Creation feature that is part of the default installation of SharePoint 2013 does not work with host-named site collections.	When you use the Self Service Site Creation feature that is part of the default installation of SharePoint 2013, you create path-based sites.
Managed paths	Managed paths for host-named site collections apply at the farm level and are available for all web applications. You have to use Windows PowerShell to create managed paths for host-named site collections.	Managed paths for path-based sites apply at the web application level. You can use Central Administration or Windows PowerShell to create managed paths for path-based site collections.

Design and manage URLs for host-named site collections

Windows PowerShell cmdlets manage URL mappings for host-named site collections and enable you to map URLs to a single site collection:

- Set-SPSiteUrl — Add or change a URL mapping for a site.
- Remove-SPSiteUrl — Remove a URL mapping from a site.
- Get-SPSiteUrl — See all URLs and associated zones for a site collection.

These cmdlets provide URL mapping functionality for host-named site collections that is similar to alternate access mapping.

Zones and host-named site collections

Host-named site collections are available through any zone. Host-named site collections are not limited to the default zone. If needed, you can implement multiple zones and use zones and host-named site collections to configure different authentication settings or policies.

You can assign up to five URLs to a single site collection by assigning one URL per zone. Even if you follow the recommended architecture by implementing only one zone, you can still assign up to five URLs to host-named site collections. This is because if a zone is not implemented by extending the web application, SharePoint 2013 uses the default zone.

For example, the following URLs could provide access to the same Internet site:

- www.Contoso.com
- www.Contoso.uk
- www.Contoso.ca
- www.Contoso.au
- www.Contoso.ie

The search crawl account requires access to content through the Default zone by using Integrated Windows authentication (NTLM or Kerberos). Because claims authentication allows multiple types of authentication in one zone, this requirement should not affect other authentication requirements.

Managed paths and host-named site collections

URLs that are configured for the same site collection may have different schemes and domains, but they must have the same managed paths, namely everything after the '/' that follows the domain must be the same. For example, http://www.Contoso.com/sites/Site1 and http://www.Fabrikam.com/sites/Site1 can both point to the same site collection but http://www.Contoso.com/sites/Site1 and http://www.bar.com/sites/Project1 cannot.

The cmdlets that manage URLs only operate on the root site collection for a host name, for example http://www.Contoso.com. These cmdlets do not operate on a managed path site collection that is underneath the root, such as http://www.Contoso.com/sites/Project1. Sites below the root of a host-named site collection will inherit the URL settings of that root host-named site collection.

Off-box termination of SSL with host-named site collections

Off-box termination of SSL occurs when a proxy server terminates an SSL request and uses HTTP to forward the request to a web server. To achieve off-box SSL termination with host-named site collections, the device that terminates the SSL connection, such as a reverse proxy server, must be capable of generating a custom HTTP header: **Front-End-Https: On**. For more information, see Use host-named site collections with off-box SSL termination, later in this article.

The protocol used for a host-named site collection depends on the value of the Url parameter that you specified when you used the Set-SPSiteURL cmdlet to map the URL to a particular zone: http or https. Ensure that the IIS bindings for the web application, SSL certificates, reverse proxy configuration, and any other configuration necessary is complete.

When to use path-based site collections

Although we recommend host-named site collections for most architectures, you should use the traditional path-based site collections and alternate access mapping if any of the following conditions apply:

- You need to use the Self Service Site Creation feature that is part of the default installation of SharePoint 2013.

  This does not apply to custom self-service site creation solutions.

- SSL termination is required but your SSL termination device cannot be configured to produce the necessary custom HTTP header.

  You can still use SSL bridging with host-named site collections with these devices if SSL termination is not a requirement.

- You plan to use different application pools for the additional security that these provide or you need to use multiple proxy groups.

  In these cases, you can use host-named site collections. However, the additional configuration that is required to map URLs for host-named site collections across multiple web applications greatly outweigh the benefits of using host-named site collections. For more information, see Use multiple web applications with host-named site collections. For more information about creating path-based site collections, see Create a site collection in SharePoint 2013.

Use host headers and host-named site collections

Host headers allows the web server to host multiple web sites on the same IP Address and Port combination. If the incoming HTTP request includes a host header name, and a matching host header is configured in IIS, IIS will respond with the content from the appropriate web site.

Host headers are configured at the Web Application (IIS website) level, they are one of the website bindings properties.

It is important to understand the distinction between Host headers in IIS and Host Named Site Collections. Host headers at the IIS website level are only intended for path-based site collections.

When using Host named site collections, SharePoint is responsible for resolving the correct site for the address based upon the incoming request passed through IIS. In most cases, applying a host header binding at the IIS website level makes it impossible to access host-named site collections through the IIS website. This is because IIS will not respond to requests for host names that differ from the host header binding.

> **Important:**
>
> If an existing web application has a host header binding set, IIS will not return pages from the host-named site collection until you remove the binding from IIS. For more information, see Update a Web application URL and IIS bindings (Office SharePoint Server 2007).

Mix host-named site collections and path-based site collections in the same web application

You can use host-named and path based site collections in the same web application. To ensure that both types of site collections are accessible to users, do not put host header bindings on the IIS website of your web application, including IIS websites for zones that are extended from the web application. If an existing web application has a host header binding set, IIS will not return pages from the host-named site collection until you remove the binding from IIS.

My Sites

When you use both types of site collections with My Sites, consider implementing your own provisioning process to create My Sites as host-named sites instead of path based sites.

Deployment and configuration for host-named site collections

Create a web application for host-named site collections

If you do not intend to configure two or more IIS websites that share the same port number on the same server, create a web application in the Default zone. Do not apply a host header binding at the IIS website level.

To create a web application for host-named site collections

1. Verify that you have the following memberships:
   - The **securityadmin** fixed server role on the SQL Server instance.
   - The **db_owner** fixed database role on all databases that are to be updated.
   - The Administrators group on the server on which you are running the Windows PowerShell cmdlet.

   An administrator can use the **Add-SPShellAdmin** cmdlet to grant permissions to use SharePoint 2013 cmdlets.

   > **Note:**
   >
   > If you do not have permissions, contact your Setup administrator or SQL Server administrator to request permissions. For additional information about Windows PowerShell permissions, see Add-SPShellAdmin.

2. On the **Start** menu, click **All Programs**.
3. Click **Microsoft SharePoint 2013 Products**.
4. Click **SharePoint 2013 Management Shell**.
5. At the Windows PowerShell command prompt (that is, PS C:\>), type the following syntax:

```
New-SPWebApplication -Name 'Contoso Sites' -port 80 -ApplicationPool ContosoAppPool -ApplicationPoolAccount (Get-SPManagedAcco
```

Create a root site collection

A root site collection is a requirement for any Web application. It is also necessary for crawling content. This site collection must have the same URL as the Web application. Currently, SharePoint prevents the creation of a host-named site collection with the same URL as a Web application. Therefore, the root site collection is created as a path-based site collection.

The following example creates an empty site collection that is the root site collection:

```
New-SPSite 'http://<servername>' -Name 'Portal' -Description 'Portal on root' -OwnerAlias 'contoso\administrator' -language 1033 -Te
```

Only the root site collection of the web application appears in the content source. Even though all other host-named site collections in the web application do not appear in the content source, by default search automatically crawls the other host-named site collections.

Create host-named site collections

You must use Windows PowerShell to create a host-named site collection. You cannot use the SharePoint 2013 Central Administration web application to create a host-named site collection, but you can use Central Administration to manage the site collection after you have created it.

You can create a host-named site collection by using the Windows PowerShell New-SPSite cmdlet with the -HostHeaderWebApplication parameter, as shown in the following example:

To create host-named site collections

1. Verify that you have the following memberships:
   - The **securityadmin** fixed server role on the SQL Server instance.
   - The **db_owner** fixed database role on all databases that are to be updated.
   - The Administrators group on the server on which you are running the Windows PowerShell cmdlet.

   An administrator can use the **Add-SPShellAdmin** cmdlet to grant permissions to use SharePoint 2013 cmdlets.

> **Note:**
>
> If you do not have permissions, contact your Setup administrator or SQL Server administrator to request permissions. For additional information about Windows PowerShell permissions, see Add-SPShellAdmin.

2. On the **Start** menu, click **All Programs**.
3. On the **Start** menu, click **All Programs**.
4. Click **Microsoft SharePoint 2013 Products**.
5. Click **SharePoint 2013 Management Shell**.
6. At the Windows PowerShell command prompt (that is, PS C:\>), type the following syntax:

```
New-SPSite 'http://portal.contoso.com' -HostHeaderWebApplication 'http://<servername>' -Name 'Portal' -Description 'Customer r
```

This creates a host-named site collection that has the URL, http://portal.contoso.com, in the SharePoint 2013 web application that has the URL ,http://webapp.contoso.com.

Use managed paths with host-named site collections

You can implement managed paths with host-named site collections. Hosters can provide multiple site collections to the same customer with each site collection sharing the customer's unique host name but differentiated by the URL path after the host name. Managed paths for host-named site collections are limited to 20 per farm. For more information, see Web application limits.

Managed paths for host-named site collections behave differently from managed paths for path-based site collections. Managed paths for host-named site collections are available to all host-named site collections within the farm regardless of the web application that the host-named site collection is in. In contrast, managed paths for path-based site collections only apply to sites within the same web application. Managed paths for path-based site collections do not apply to path-based site collections in other web applications. Managed paths for one type of site collection do not apply to the other type of site collection.

To create a managed path, you must first create a site collection with the desired base URL. For example, to create http://teams.contoso.com/*finance* you must first create the site collection for http://teams.contoso.com.

To create a managed path for use with host-named site collections, use the Windows PowerShell **New-SPManagedPath** cmdlet with the **-HostHeader** parameter, as shown in the following example:

```
New-SPManagedPath 'departments' -HostHeader
```

You can also use the **–explicit** parameter to create explicit managed paths.

The following example shows a host-named site collection created at a managed path:

```
New-SPSite 'http://portal.contoso.com/departments/marketing' -HostHeaderWebApplication 'http://<servername>' -Name 'Marketing' -Desc
```

To remove an existing managed path, use the Windows PowerShell **Remove -SPManagedPath** cmdlet, as shown in the following example:

```
Remove-SPManagedPath 'departments' -HostHeader
```

You can use Windows PowerShell to remove a managed path even if a site collection exists. If you remove a managed path, the site collection can no longer be accessed. To access the existing site collection, use Windows PowerShell to recreate the managed path.

Map URLs to host-named site collections

When you create a new hostname site collection, the default alternate access mappings will still exist but cannot be used. Use Windows PowerShell commands to manage URL mappings for host-named site collections.

Add a mapping to an existing site:

```
Set-SPSiteUrl (Get-SPSite 'http://teams.contoso.com') -Url 'http://teamsites.contoso.com' -Zone Intranet
```

Each URL mapping is applied to a single zone. Use one of the following zone names when you map URLs:

- Default
- Intranet
- Internet
- Custom
- Extranet

If you do not specify the **-Zone** parameter and the URL mapping entry is new, the default zone is used.

Remove a mapping for a site:

```
Remove-SPSiteUrl 'http://teamsites.contoso.com'
```

Display all URL mappings for a site:

```
Get-SPSiteUrl -Identity (Get-SPSite 'http://teams.contoso.com')
```

Configure SSL certificates for host-named site collections

You can configure a single web application that uses SSL and then create multiple host-named site collections within that web application. To browse to a site over SSL, you have to install and assign a server certificate to the IIS website. Each host-named site collection in a web application will share the single server certificate that you assigned to the IIS website.

You need to acquire a wildcard certificate or subject alternate name (SAN) certificate and then use a host-named site collection URL format that matches that certificate. For example, if you acquire a *.contoso.com wildcard certificate, you must generate host-named site collection URLs such as https://site1.contoso.com, https://site2.contoso.com, and so on, to enable these sites to pass browser SSL validation. However, if you require unique second-level domain names for sites, you must create multiple web applications rather than multiple host-named site collections.

To configure SSL for host-named site collections, enable SSL when you create the web application. This will create an IIS website with an SSL binding instead of an HTTP binding. After you create the web application, open IIS Manager and assign a certificate to that SSL binding. You can then create site collections in that web application.

If you are implementing multiple zones with host-named site collections, ensure that the configuration of certificates and bindings (SSL or HTTP) is appropriate for each zone and corresponding IIS site.

Use host-named site collections with off-box SSL termination

You can use host-named site collections with off-box SSL termination. There are several requirements to use SSL termination with host-named site collections:

- At least one IIS site should have a binding on port 80 (or whatever port the terminator forwards the request to). Microsoft recommends that you use the IIS site of a web application (or the IIS site of a zone for a web application) with HTTP/80.
- The SSL terminator or reverse proxy must preserve the original HTTP host header from the client.
- If the client SSL request is sent to the default SSL port (443), then the SSL terminator or reverse proxy must forward the decrypted HTTP request to the front-end web server on the default HTTP port (80). If the client SSL request is sent to a non-default SSL port, then the SSL terminator or reverse proxy must forward the decrypted HTTP request to the front-end web server on the same non-default port.
- The device that terminates the SSL connection, such as a reverse proxy server, must be capable of generating a custom HTTP header: **Front-End-Https: On**. This is the same custom header that Outlook Web Access (OWA) uses: **Front-End-Https: On/Off**. More information about this custom header is included later in this section.

To use host-named site collections with off-box SSL termination, configure your web application as you normally would for SSL termination and ensure that it meets the requirements described above. In this scenario, SharePoint 2013 will use HTTPS instead HTTP to render links of its host-named site collections in that web application.

Reverse proxy servers can publish SharePoint 2013 host-named site collections and perform off-box SSL termination. In this scenario, the reverse proxy server changes the connection type between the end user and the SharePoint web front-end server from SSL/TLS to HTTP or vice versa. Reverse proxy servers in this scenario must insert an additional HTTP header in the user's request when it forwards the request to the SharePoint web front-end server. This additional HTTP header indicates to SharePoint 2013the type of connection the end user initiated so that SharePoint 2013 renders URLs appropriately in its response. The HTTP header name is "Front-End-Https" and its acceptable values are as follows.

Table: Front-End-Https header values

Value	Description
On	The reverse proxy server received the end user's request over an encrypted (SSL or TLS) HTTPS connection. For example, **Front-End-Https: On**.
Off	The reverse proxy server received the end user's request over an unencrypted HTTP connection.

The values are not case-sensitive. For example, on, ON, On, and oN are acceptable.

This custom header works only with host-named site collections. It does not work with path-based site collections.

Enable apps in environments with multiple zones

The March 2013 Public Update enables you to configure an app domain for each web application zone and use alternate access mapping and host-header web application configuration. Before the release of this update, you could only host one app domain and it had to be in the Default zone. You could not use the app domain on alternate access mappings or host-header web application configurations.

To resolve this issue, apply the SharePoint 2013 Cumulative Update Server Hotfix Package: March 12, 2013, see Updates for SharePoint 2013.

Migrate path-based site collections to host-named site collections

Determine host-named site collections in existing web applications

When you migrate from SharePoint 2010 to SharePoint 2013, we recommend that you determine how SharePoint 2010 sites were created. If sites were created as path-based sites, consider migrating these sites to host-named site collections. If host-named and path-based sites were implemented together, identify the sites that were created as path-based sites and consider migrating these sites to host-named site collections. To accomplish this, look for the 'HostHeaderIsSiteName' flag.

The following example determines if a site within a given web application is created as host-named or path based:

```
$webApp = Get-SPWebapplication 'http://portal.contoso.com'

foreach($spSite in $webApp.Sites)
{
if ($spSite.HostHeaderIsSiteName)
{ Write-Host $spSite.Url 'is host-named' }
else
{ Write-Host $spSite.Url 'is path based' }
}
```

Convert path-based site collections to host-named site collections

You can convert path-based site collections to host-named site collections and host-named site collections to path-based site collections. You have to use the Windows PowerShell backup and restore cmdlets to convert site collections. You cannot use the the SharePoint Central Administration website or Windows PowerShell cmdlets that attach and detach, or mount and dismount content databases to convert site collections.

The following example converts a standard site collection to a host-named site collection:

```
Backup-SPSite -Identity 'http://portalOld.contoso.com' -Path 'c:\Backup\portalContoso.bak' -Force -UseSQLSnapShot
Restore-SPSite -Identity 'http://portal.contoso.com' -Path 'c:\Backup\portalContoso.bak' -DatabaseName 'portal_content' -Force -Host
```

Use multiple web applications with host-named site collections

If you use more than one web application, you add more operational overhead and complexity to the system. We recommend that you use one web application for site collections. However, the following reasons might influence you to implement site collections across multiple web applications:

- An organization's security policies require separate web applications or application pools.
- Web applications need to be configured differently.
- An organization requires the use of multiple proxy groups.

It is more complex to implement host-named site collections with multiple web applications in a farm because you must complete more configuration steps. For example, URLs with host-named sites might be spread across multiple web applications that share the same port in a single farm. This scenario requires more configuration steps to ensure that requests are mapped to the correct web applications. You have to manually configure the mappings on each web server in the farm by configuring a separate IP address to represent each web application. You also have to create and manage host-header bindings to assign unique IP addresses for each site. Scripts can manage and replicate this configuration across servers; however, this adds complexity to the solution. Each unique URL also requires a mapping in DNS. Generally speaking, if multiple web applications are a requirement, we recommend path-based site collections with alternate access mapping.

The following two tables contrast three different design choices to implement site collections. These tables are intended to help you understand the consequences of each approach and how configuration varies depending on the architecture.

Table: Results of different design choices to provision site collections

	Host-named site collections with all sites in a farm consolidated in one web application	Path-based site collections with alternate access mapping and multiple web applications	Host-named site collections with multiple web applications in a farm
Site collection provisioning	Use Windows PowerShell or a custom site collection provisioning solution to provision sites.	Use Central Administration or Windows PowerShell to deploy sites.	Use Windows PowerShell or a custom site collection provisioning solution to provision sites.
URL management	You can map all site collections in DNS to point to a single IP address which represents the web application.	If you implemented more than one zone, you configure alternate access mapping for each site URL. Each zone also requires a mapping in DNS.	Additional configuration is required to ensure that requests for sites that share the same port are mapped to the correct web application. Each unique host name also requires a mapping in DNS. This configuration is manual and you must complete it on each web server in a farm for each site.
Additional URLs	You can assign up to five URLs to a host-named site collection, one per zone. It is not necessary to extend the web application to multiple zones. If a zone is not implemented the default zone is used.	The number of URLs for a site collection is limited to five because this is the number of zones that is allowed.	You can assign up to five URLs to a host-named site collection, one per zone. It is not necessary to extend the web application to multiple zones. If a zone is not implemented the default zone is used.
Service applications	All sites in the farm use a single service application group.	You can implement custom service application groups for different web applications.	You can implement custom service application groups for different web applications.
Zones	You do not have to implement multiple zones to implement different URLs for the same site collection. If a zone is not implemented the default zone is used.	Zones are necessary to implement different URLs for the same site collection.	You do not have to implement multiple zones to implement different URLs for the same site collection. If a zone is not implemented the default zone is used.
Authentication	With one web application, authentication options are limited to five zones. However, you can implement many authentication methods on one zone.	You can implement different authentication and zone designs for each web application.	You can implement different authentication and zone designs for each web application.
Authentication	Provides client scripting isolation between domain URLs.	You can isolate web applications into dedicated application pools, if desired, to achieve process isolation. Provides isolation between domain URLs.	You can isolate web applications into dedicated application pools, if desired, to achieve process isolation. Provides isolation between domain URLs.
Policy	You can use zones to assign different policies to host-named sites.	You can use policies at the web application level to enforce permissions, regardless of permissions that are configured on individual sites or documents. Additionally, you can implement different policies for different zones.	You can implement different policies for different web applications to enforce permissions, regardless of permissions that are configured on individual sites or documents. Additionally, you can implement different policies for different zones.

Scalability numbers that might also affect design decisions include the recommended maximums for site collections, content databases, and managed paths.

The following table summarizes the configuration that is necessary to manage URLs based on each of the three design options presented in this article.

Table: Configuration required for different site collection designs

	Host-named site collections with all sites in a farm consolidated into one web application	Path-based site collections with alternate access mapping and multiple web applications	Host-named site collections with multiple web applications in a farm
Within SharePoint Server	• Create the web application. • Create a root-site collection that is not accessible to users (for example, https://HNSC01.fabrikam.com). • Create the host-named site collections with the host header (for example, https://intranet.fabrikam.com). • Optionally add more URLs for each site collection and configure zones by using Set-SPSiteUrl. (In corporate portal design samples there's no need because there's only one zone.)	• Create the web application with the host header (for example, https://intranet.fabrikam.com). • Optionally configure alternate access mapping. In the design sample there's no need because there's only one zone). • Create the root path-based site collection.	• Create the web application. • Create a root-site collection that is not accessible to users (for example, https://HNSC01.fabrikam.com). • Create the host-named site collections with the host header (for example, https://intranet.fabrikam.com). • Optionally add more URLs for each site collection and configure zones by using Set-SPSiteUrl. (In corporate portal design samples there's no need because there's only one zone.)
Within IIS	Associate an SSL certificate (wildcard certificate or SAN certificate) for all host-named site (domain) in the web application.	Associate an SSL certificate in IIS for each zone (each zone is a separate web application in IIS).	Associate an SSL certificate (wildcard certificate or SAN certificate) for a host-named site (domain) in the web applications.

			On each web server in the farm and for each web application that shares a port: • Configure a separate IP address to represent each web application. • Edit the IIS web site binding manually to remove the host header binding that was created when the web application was created and replace this with an IP address binding.

If you use multiple web applications on different IP addresses, you might need to complete additional configuration for the NIC, DNS, and the load balancer for each server.

Create multiple web applications with host-named site collections

To run multiple web applications on the same server and port in combination with host-named site collections, you have to assign different IP addresses to the web applications. This type of architecture requires that you add IP addresses to the web servers and configure the network router to point host names to the IP address of its web application.

> **Note:**
>
> You can create a web application that does not have a host header. If you create a web application that does not have a host header, you cannot create multiple web applications with host-named site collections on the same web server.

The process that creates multiple web applications for a host-named site collections includes the following tasks:

- Create the multiple web applications.
- Add a new virtual IP address in IIS on each web server in the farm.

Create multiple web applications for host-named site collections

The following example creates a web application:

```
New-SPWebApplication -Name 'webapp' -hostHeader 'webapp.contoso.com' -port 80 -ApplicationPool ContosoAppPool -ApplicationPoolAccour
```

Repeat this task for each web application.

Add virtual IP addresses in IIS

The IP bindings must be applied on all servers that will host the web application. Set the sleep command to 60 seconds to make sure that the IP bindings are set to all servers in the farm before the existing host header on the web application is removed. Remote scripting can be used for this job.

Use the following commands to add unique IP bindings to each of the web applications that you created and then remove the host header binding from these web applications.

```
Import-Module WebAdministration
# add empty binding to webapp on IP 192.168.10.20
New-WebBinding -Name 'webapp' -IPAddress '192.168.10.20' -HostHeader ''
Sleep 60
# remove existing binding webapp.contoso.com from existing web application
Get-WebBinding -Name 'webapp' -HostHeader 'webapp.contoso.com' | Remove-WebBinding
```

See also

Get-SPSiteUrl
Set-SPSiteUrl
Remove-SPSiteUrl

Plan logical architectures for SharePoint 2013

Configure alternate access mappings for SharePoint 2013

SharePoint 2013

Applies to: *SharePoint Server 2013, SharePoint Foundation 2013*

Topic Last Modified: *2013-12-18*

Summary: Learn how to configure alternate access mappings in SharePoint 2013.

Each web application can be associated with a collection of mappings between internal and public URLs. Both internal and public URLs consist of the protocol and domain portions of the full URL (for example, https://www.fabrikam.com). A public URL is what users type to access the SharePoint site, and that URL is what appears in the links on the pages. Internal URLs are in the URL requests that are sent to the SharePoint site. Many internal URLs can be associated with a single public URL in multi-server farms (for example, when a load balancer routes requests to specific IP addresses to various servers in the load-balancing cluster).

Each web application supports five collections of mappings per URL. The five collections correspond to five zones (default, intranet, extranet, Internet, and custom). When the web application receives a request for an internal URL in a particular zone, links on the pages returned to the user have the public URL for that zone. For more information, see Plan alternate access mappings for SharePoint 2013.

Manage alternate access mappings

1. On the SharePoint Central Administration website, click **System Settings**.
2. On the **System Settings** page, in the **Farm Management** section, click **Configure alternate access mappings**.

For information about how to perform this procedure using a Windows PowerShell cmdlet, see New-SPAlternateUrl.

Add an internal URL

1. On the **Alternate Access Mappings** page, click **Add Internal URLs**.
2. If the mapping collection that you want to change is not specified, then choose one. In the **Alternate Access Mapping Collection** section, on the **Alternate Access Mapping Collection** menu, click **Change alternate access mapping collection**.
3. On the **Select an Alternate Access Mapping Collection** page, click a mapping collection.
4. In the **Add internal URL** section, in the **URL protocol, host and port** box, type the new internal URL (for example, https://www.fabrikam.com).
5. In the **Zone** list, click the zone for the internal URL.
6. Click **Save**.

Edit or delete an internal URL

> **Note:**
>
> You can't delete the last internal URL for the default zone.

1. On the **Alternate Access Mappings** page, click the internal URL that you want to edit or delete.
2. In the **Edit internal URL** section, change the URL in the **URL protocol, host and port** box.
3. In the **Zone** list, click the zone for the internal URL.
4. Do one of the following:
   - Click **Save** to save your changes.
   - Click **Cancel** to discard your changes and return to the **Alternate Access Mappings** page.
5. Click **Delete** to delete the internal URL.

For information about how to perform this procedure using a Windows PowerShell cmdlet, see Remove-SPAlternateUrl.

Edit public URLs

> **Note:**
>
> There must always be a public URL for the default zone.

1. On the **Alternate Access Mappings** page, click **Edit Public URLs**.

2. If the mapping collection that you want to change is not specified, then choose one. In the **Alternate Access Mapping Collection** section, on the **Alternate Access Mapping Collection** menu, click **Change alternate access mapping collection**.
3. On the **Select an Alternate Access Mapping Collection** page, click a mapping collection.
4. In the **Public URLs** section, you can add new URLs or edit existing URLs in any of the following text boxes:
    - **Default**
    - **Intranet**
    - **Extranet**
    - **Internet**
    - **Custom**
5. Click **Save**.

Map to an external resource

You can also define mappings for resources outside intranet applications. To do so, you must supply a unique name, initial URL, and a zone for that URL. (The URL must be unique to the farm.)

1. On the **Alternate Access Mappings** page, click **Map to External Resource**.
2. On the **Create External Resource Mapping** page, in the **Resource Name** box, type a unique name.
3. In the **URL protocol, host and port** box, type the initial URL.
4. Click **Save**.

Plan alternate access mappings for SharePoint 2013

SharePoint 2013

Applies to: *SharePoint Server 2013, SharePoint Foundation 2013*

Topic Last Modified: *2013-12-18*

Summary: Learn how to plan for alternate access mappings in SharePoint 2013.

Alternate access mappings direct users to the correct URLs during their interaction with SharePoint 2013 (while browsing to the home page of a SharePoint 2013 website, for example). Alternate access mappings enable SharePoint 2013 to map web requests to the correct web applications and sites, and they enable SharePoint 2013 to serve the correct content back to the user.

Because the features of alternate access mapping are deprecated, we recommend that you use host-named site collections over alternate access mappings.

For additional information about how to plan for host-named site collections, see Host-named site collection architecture and deployment (SharePoint 2013).

In this article:

Alternate access mappings were implemented because there are common Internet deployment scenarios in which the URL of a web request received by Internet Information Services (IIS) differs from the URL that was typed by a user. This is most likely to occur in deployment scenarios that include reverse proxy publishing and load balancing.

> **Note:**
>
> Alternate access mappings must be configured for load balancing, even though it generally does not apply to host header site collections. The default zone public URL should be set to a domain URL that is appropriate for all users to see. Unless you do this, the names of web servers or their IP addresses might be displayed in parameters that were passed between pages within SharePoint 2013.

About alternate access mappings

Alternate access mappings enable a web application that receives a request for an internal URL in one of the five zones to return pages that contain links to the public URL for the zone. You can associate a web application by using a collection of mappings between internal and public URLs. *Internal* refers to the URL of a web request as it is received by SharePoint 2013. *Public* refers to the URL by which SharePoint will format links that correspond to requests that match one of the internal URLs on that zone when it returns a response. The public URL is the base URL that SharePoint 2013 uses in the pages that it returns. If the internal URL was changed by a reverse proxy device, it can differ from the public URL.

> **Note:**
>
> Host-named site collections can't use alternate access mappings. Host-named site collections are automatically considered in the Default zone, and the URL of the request must not be changed between the user and the server.

Multiple internal URLs can be associated with a single public URL. Mapping collections can contain up to five authentication zones. But each zone can have only a single public URL. Mapping collections correspond to the following authentication zones:

- Default
- Intranet
- Internet
- Custom
- Extranet

Reverse proxy publishing

A reverse proxy is a device that sits between users and your web server. All requests to your web server are first received by the reverse proxy device and, if those requests pass the proxy's security filtering, the proxy forwards the requests to your web server.

Alternate access mapping integration with authentication providers

Alternate access mappings allow you to expose a web application in as many as five different zones, with a different IIS website backing each zone.

> **Note:**
>
> Some people mistakenly refer to this as having up to five different web applications sharing the same content databases. In reality, there is just one web application.

Not only do these zones allow you to use multiple URLs to access the same web application, they also allow you to use multiple authentication providers to access the same web application.

When extending a web application into a zone, you have to use Windows authentication provided by IIS. After the web application has extended into the zone, you can change the zone to use a different type of authentication.

Use the following procedure to change the authentication configuration for a zone.

To change the authentication type for a zone

1. From **Administrative Tools**, open Central Administration.
2. On the Central Administration home page, click **Application Management**.
3. On the **Application Management** page, in the **Application Security** section, click **Authentication providers**.
4. On the **Authentication Providers** page, select your web application, which is listed in the **Web Application** box.
5. Click the name of the zone whose authentication configuration you want to change.

   > **Note:**
   >
   > You'll be able to select only from among zones that have a backing IIS website. These zones were assigned an IIS website during the "Extend an existing web application" procedure.

6. On the **Edit Authentication** page, in the **Claims Authentication Types** section, select the authentication type that you want to use for this zone:
   - **Windows authentication (default value)**
   - **Basic authentication**
   - **Forms based authentication (FBA)**
   - **Trusted Identity provider**
7. Change any other authentication configuration settings that you want to change, and click **Save**.

At this point, you can change authentication configuration settings for any other zone. You can configure completely independent authentication settings for different zones accessing the same content. For example, you might configure some content to be anonymously available while other content requires credentials. You could configure one zone to have anonymous access enabled and all other forms of authentication disabled, guaranteeing that only the anonymous content will be available. At the same time, another zone can have anonymous access disabled while NTLM authentication is enabled, guaranteeing that only authenticated access will be enabled. In addition, you can have different types of accounts to access the same content: one zone can be configured to use Active Directory accounts in Windows while another zone can be configured to use non-Active Directory accounts that use ASP.NET forms-based authentication.

Alternate access mapping integration with web application policies

Web application policies allow administrators to grant or deny access to accounts and security groups for all sites exposed through a zone. This can be useful for many scenarios.

For example, the SharePoint 2013 search crawler must undergo the same authorization infrastructure as any other user: it can only crawl content that it has access to. But users would still like search to crawl restricted content so that authorized users can find that content in search results. The search service uses a Full Read policy on the web applications to give its crawler permission to read all content on that web application. That way, it can crawl and index all existing and future content, even content to which the site administrator had not explicitly given it access.

Another example would be helpdesk personnel who need administrative access to SharePoint 2013 sites so that they can help users. To do this, you can create a web application policy that grants the helpdesk staff accounts Full Control permission so that they have full administrative access to all current and future sites on the web application.

Because policies are tied to both web applications and their zones, you can guarantee that the policy that you have applied to one zone does not affect other zones. This can be useful if you have content exposed both on the corporate network and to the Internet. For example, suppose that you have given a helpdesk staff account Full Control permission over a web application's zone that is assigned to the corporate network. If someone were to try to use that account to access the site over the Internet, that Full Control policy wouldn't apply because it would recognize that the URL is in a different zone. Therefore, the account wouldn't automatically be given administrative access to the site.

Alternate access mapping and external resource mapping

SharePoint 2013 allows you to extend the alternate access mapping functionality to content that is not hosted within the SharePoint 2013 farm. To configure this functionality, browse to the **Alternate Access Mappings** page, and then click **Map to External Resource**. You'll then be asked to create an entry for an external resource, which you can think of as another web application. After you have an external resource, you can assign different URLs and zones to it in the same manner that you do for web applications. This feature is not used in SharePoint 2013, but third-party products that build onto SharePoint 2013 can use it.

For example, the search technology in SharePoint 2013 can crawl content external to the farm, such as file shares and websites. If that content is available at different URLs on different networks, you would want search to return results by using the appropriate URLs for the user's current network. By using alternate access mapping's

external resource mapping technology, search can remap the external URLs in its results to match the user's zone.

See also

Configure alternate access mappings for SharePoint 2013

Update a web application URL and IIS bindings for SharePoint 2013

SharePoint 2013

Applies to: *SharePoint Server 2013, SharePoint Foundation 2013*

Topic Last Modified: *2014-06-12*

Summary: Learn how to update a web application and IIS bindings for SharePoint 2013.

After you have extended a web application into a zone with a set of Internet Information Services (IIS) bindings and alternate access mapping URLs, you might decide that you want to use a different URL to reach the web application. For example, you might have originally created a web application to use HTTP and later decide to use SSL instead. Or, you might create a web application to use the www.contoso1.com host header and then decide to use the www.contoso2.com host header instead. This article provides detailed guidance for changing the URL and IIS bindings of a web application.

> **Note:**
>
> This article focuses on changing the existing URL and IIS bindings of a web application. If you want to add additional URLs and IIS bindings to a web application, you can do so by extending the web application into an unused zone.

In this article:

About updating a web application URL and IIS bindings

Unlike typical IIS applications, you cannot simply use IIS Manager or other IIS metabase tools to modify the bindings of IIS web applications that have been extended with SharePoint 2013.

If you modify the IIS bindings of a web application by adding a host header binding or SSL port or by changing a port number, SharePoint 2013 will not be aware of these changes and will not update the web application's alternate access mapping URLs. If you update the web application's alternate access mappings to change a host header, switch to an SSL URL, or change a port number, SharePoint 2013 will not automatically update your IIS bindings to match.

To update the URL or IIS bindings of a web application, unextend and reextend the web application and reconfigure the alternate access mapping URLs or IIS website bindings.

We do not recommend reusing the same IIS website for your HTTP and SSL hosting. Instead, extend a dedicated HTTP and a dedicated SSL website, with each assigned to its own alternate access mapping zone and URLs.

For more information about alternate access mappings, see Plan alternate access mappings for SharePoint 2013.

Alternate access mapping collections correspond to the following authentication zones:

- Default
- Intranet
- Internet
- Custom
- Extranet

Unextending and reextending a web application

If you need to change your IIS bindings, unextend the web application from the zone that the web application has been extended into (without deleting the web application), and then reextend the web application into the same zone. Consider trying these migration procedures in a test environment before deploying them in a production environment.

Unextend the web application from the zone by using the **Remove SharePoint from IIS Web site** link on the **Central Administration Application Management** page, as described in the following procedure.

To unextend a web application

1. On the SharePoint Central Administration website, on the **Application Management** page, in the **Web Applications** section, click **Manage web applications**.
2. On the **Web Applications** menu, click the web application you want to unextend, click **Delete**, and then click **Remove SharePoint from IIS Web Site**.
3. On the **Remove SharePoint From Web Site** page, click the web application you want to unextend.
4. In the **Select IIS Web site and zone to remove** list, click the IIS website and zone you want to remove. Because a web application can be extended in up to five

zones, make sure you select the correct IIS website and zone.

5. In the **Delete IIS Web sites** section, click **Yes** if the IIS website is hosting only SharePoint Products and Technologies content. If the IIS web site is hosting other content, you might not want to delete the web site. In that case, click **No**.
6. Click **OK**. This action does not delete the web application, nor does it delete the content databases of the web application.

After you have unextended the web application, you can reextend the web application to the same zone by using your updated bindings.

To reextend a web application

1. On the SharePoint Central Administration website, on the **Application Management** page, in the **Web Applications** section, click **Manage web applications**.
2. Click the web application you want to extend. On the ribbon, click **Extend**.
3. In the **IIS Web Site** section, if you have already created an IIS website with the appropriate bindings for SharePoint 2013 to use, click the **Use an existing IIS Web site** option and select the IIS website from the list. Otherwise, click the **Create a new IIS Web site** option.
4. In the **Port**, **Host Header**, and **Use Secure Sockets Layer (SSL)** fields, type the IIS bindings you want to use.
5. In the **Load Balanced URL** section, in the **URL** field, type the URL that users will use to locate this web application. If you are using a load balancer or reverse proxy, this is the URL of the load balancer or reverse proxy.
6. In the **Load Balanced URL** section, in the **Zone** list, click the zone that you previously selected.
7. Click **OK**.

Additional steps for updating a web application URL and IIS bindings

To complete the process of updating a web application URL or IIS bindings, perform the additional steps listed in this section after you have reextended the web application into the same zone.

Update the alternate access mapping URLs for the zone

If you are using a load balancer or a reverse proxy, make sure that your internal URLs are updated in the alternate access mappings to reflect the new IIS bindings. In addition, update your load balancer rules or your reverse proxy rules to align with the new IIS bindings.

Apply an SSL certificate

If the new IIS bindings use SSL, apply an SSL certificate to the new IIS website assigned to your zone. For more information about how to configure SSL, see Configuring Secure Sockets Layer (IIS 6.0).

Apply an authentication method

When you reextend your web application, the default SharePoint 2013 authentication method for the zone is Integrated Windows authentication. If you want to use an authentication method other than Integrated Windows, explicitly apply the authentication method you want to use. For more information about authentication methods, see Plan for user authentication methods in SharePoint 2013.

For more information about how to configure authentication for SharePoint Server 2013, see Configure authentication infrastructure in SharePoint 2013.

Verify the Search start addresses and My Site settings

Verify that the SharePoint 2013 Search start addresses are correct for your content sources. If the SharePoint 2013 Search start addresses are incorrect, update them. Also, verify that your My Site settings, including Personal Search Center, Personal Site Provider, and default RSS feed, are correct. If your My Site settings are incorrect, update them. For more information about how to configure My Site settings, see Configure My Sites in SharePoint Server 2013.

For more information about Search start addresses, see Plan crawling and federation in SharePoint Server 2013.

Verify the trusted file locations

If Excel Services in SharePoint Server 2013 is part of your deployment, verify that your trusted file locations are configured correctly. If your trusted file locations are configured incorrectly, update them. For more information about trusted file locations, see Configure Excel Services in SharePoint Server 2013.

Redeploy solutions

When you remove SharePoint 2013 from an IIS website, if you are removing the last (or only) website that is associated with the web application, any web application solutions you have deployed will also be removed. If you need these solutions, redeploy them. For additional information about how to manage solutions, see Install and manage solutions for SharePoint 2013