# Planning guide for server farms and environments for Microsoft SharePoint Server 2010

Microsoft Corporation

Published: January 2011

Author: Microsoft Office System and Servers Team (itspdocs@microsoft.com)

## Abstract

This book provides information and guidelines for making decisions about system architecture for a deployment of Microsoft SharePoint Server 2010. Subjects include system requirements, authentication, and business continuity management. Capacity planning information is provided in a separate book (link follows). The audiences for this book are business application specialists, line-of-business specialists, information architects, IT generalists, program managers, and infrastructure specialists who are planning a solution based on SharePoint Server 2010. This book is part of a set of four planning guides that provide comprehensive IT planning information for SharePoint Server.

For information about planning for capacity and performance in SharePoint Server 2010, see Capacity planning for Microsoft SharePoint Server 2010 (http://go.microsoft.com/fwlink/?LinkID=208221).

For information about planning for sites and solutions created by using SharePoint Server, see Planning guide for sites and solutions for Microsoft SharePoint Server 2010, Part 1 (http://go.microsoft.com/fwlink/?LinkID=196150) and Planning guide for sites and solutions for Microsoft SharePoint Server 2010, Part 2 (http://go.microsoft.com/fwlink/?LinkId=208024).

The content in this book is a copy of selected content in the SharePoint Server 2010 technical library (http://go.microsoft.com/fwlink/?LinkId=181463) as of the publication date. For the most current content, see the technical library on the Web.

# Contents

Getting help ........ 13

Technical diagrams (SharePoint Server 2010) ........ 14
- Models ........ 14
- Tips for printing posters ........ 26

Plan for server farms and environments (SharePoint Server 2010) ........ 27

System requirements (SharePoint Server 2010) ........ 28

Hardware and software requirements (SharePoint Server 2010) ........ 29
- Overview ........ 29
- Hardware requirements—Web servers, application servers, and single server installations ........ 29
- Hardware requirements—Database servers ........ 30
- Software requirements ........ 31
  - Minimum requirements ........ 31
  - Optional software ........ 35
- Access to applicable software ........ 36

Plan browser support (SharePoint Server 2010) ........ 39
- About planning browser support ........ 39
- Key planning phase of browser support ........ 39
  - Browser support levels ........ 39
  - Browser support matrix ........ 40
  - Browser details ........ 41
  - Browser compatibility for publishing sites ........ 57
- ActiveX controls ........ 57

URL path length restrictions (SharePoint Server 2010) ........ 58
- Understanding URL and path lengths ........ 58
  - SharePoint URL composition ........ 58
  - URL Encoding ........ 59
  - URL parameters ........ 60
- URL path length limitations ........ 61
  - SharePoint URL path length limitations ........ 61
  - Internet Explorer URL length limitations ........ 62
- Resolving URL length problems ........ 62

IP support (SharePoint Server 2010) ........ 63

Windows Server 2008 R2 and SharePoint Server 2010: Better Together (white paper) ........ 65

SQL Server 2008 R2 and SharePoint 2010 Products: Better Together (white paper) (SharePoint Server 2010) ........ 66

Business Productivity at Its Best: Microsoft Office 2010 and SharePoint Server 2010 Better Together (white paper) ........ 67

Logical architecture planning (SharePoint Server 2010) ........ 68

Services architecture planning (SharePoint Server 2010) ........ 69
About service applications ........ 70
Services infrastructure and design principles ........ 73
Deploying services ........ 73
More granular configuration of services ........ 73
Service application groups ........ 73
Logical architecture ........ 74
Connections for service applications ........ 76
Service application administration ........ 76
Deploying service applications across farms ........ 76
Design guidance ........ 77
Deploying cross-farm services ........ 77
Planning considerations for services that access external data sources ........ 78
Example architectures ........ 79
Single farm, single service application group ........ 79
Advantages ........ 79
Disadvantages ........ 80
Recommendations ........ 80
Single farm, multiple service application groups ........ 80
Connecting to multiple Managed Metadata service applications ........ 82
Advantages ........ 83
Disadvantages ........ 83
Recommendations ........ 83
Enterprise services farms ........ 84
Published content–only farms (all service applications are remote) ........ 85
Collaboration farms (mix of local and remote service applications) ........ 85
Farms for specialized departments (mix of local and remote service applications) ........ 86
Specialized service farms ........ 86
Cross-organization farms ........ 87

Logical architecture components (SharePoint Server 2010) ........ 89
Server farms ........ 89
Service applications ........ 90
Capacity ........ 90
Sharing and isolation ........ 90

Configurable items ........ 91
Administration ........ 92
Planning recommendations ........ 92
Application pools ........ 93
Capacity ........ 93
Sharing and isolation ........ 93
Configurable items ........ 93
Administration ........ 93
Planning recommendations ........ 93
Web applications ........ 94
Sharing and isolation ........ 94
Configurable items ........ 94
Administration ........ 95
Planning recommendations ........ 95
Zones ........ 95
Capacity ........ 95
Sharing and isolation ........ 95
Configurable items ........ 96
Administration ........ 97
Planning recommendations ........ 97
Policy for a Web application ........ 97
Capacity ........ 98
Sharing and isolation ........ 98
Configurable items ........ 98
Administration ........ 99
Planning recommendations ........ 99
Content databases ........ 99
Capacity ........ 99
Sharing and isolation ........ 99
Configurable items ........ 100
Administration ........ 100
Planning recommendations ........ 101
Site collections ........ 101
Capacity ........ 101
Sharing and isolation ........ 101
Configurable items ........ 102
Administration ........ 103
Planning recommendations ........ 103
Sites ........ 104
Capacity ........ 104
Sharing and isolation ........ 104
Configurable elements ........ 104

Administration ........ 104
Host-named site collections ........ 105
Capacity ........ 105
Sharing and isolation ........ 105
Administration ........ 105
My Sites ........ 105

Design sample: Corporate deployment (SharePoint Server 2010) ........ 106
About the design samples ........ 107
Company Internet site ........ 108
Overall design goals ........ 109
Server farms ........ 110
Users, zones, and authentication ........ 111
Classic mode authentication design sample ........ 112
Zones ........ 114
Services ........ 116
Authoring and publishing alternatives ........ 117
Administration sites ........ 118
Application pools ........ 118
Web applications ........ 119
Site collections ........ 119
Content databases ........ 123
Zones and URLs ........ 125
Designing load-balanced URLs ........ 126
Using explicit and wildcard inclusions for URL paths ........ 127
Zone policies ........ 130

Plan for host-named site collections (SharePoint Server 2010) ........ 131
About host-named site collections ........ 131
About host headers ........ 132
Create a host-named site collection ........ 133
Programmatically create a host-named site collection ........ 133
Use managed paths with host-named site collections ........ 134
Expose host-named sites over HTTP or SSL ........ 134
Configure SSL for host-named site collections ........ 134
Use host-named site collections with off-box SSL termination ........ 135

Hosted environments (SharePoint Server 2010) ........ 136

Model: Hosting architectures for SharePoint Server 2010 ........ 137

White paper: SharePoint 2010 for hosters (SharePoint Server 2010) ........ 138

Virtualization planning (SharePoint Server 2010) ........ 139

Virtualization support and licensing (SharePoint Server 2010) ........ 140
- SharePoint 2010 Products support for virtualization ........ 140
- Server virtualization using Hyper-V technology ........ 140
- Operating system environment (OSE) licensing ........ 140
- SharePoint 2010 Products licensing ........ 141

Hyper-V virtualization requirements (SharePoint Server 2010) ........ 142
- Hardware ........ 142
- Software ........ 142

Plan virtual architectures (SharePoint Server 2010) ........ 144
- Virtual versus physical architectures ........ 144
  - Virtualizing Web servers and application servers ........ 144
  - Virtualizing SQL Server ........ 144
  - Virtualizing other servers in the environment ........ 145
    - Active Directory ........ 145
    - Gateway products ........ 145
  - Testing side by side ........ 146
- Example virtual architectures for small to medium size farms ........ 147
- Example virtual architectures for medium to large size farms ........ 149

Plan for virtualization (SharePoint Server 2010) ........ 153
- Create a plan for deploying SharePoint Server 2010 in a virtual environment ........ 153

Capacity management and high availability in a virtual environment (SharePoint Server 2010) ........ 158
- Virtualization overview ........ 158
- Capacity management ........ 159
- Virtualization server capacity and sizing ........ 159
- Creating and refining the architectures ........ 160
  - Create the architecture ........ 161
  - Analyze the architectures ........ 161
  - Refine the architecture ........ 164
- Additional options for improving the architecture ........ 168

Plan authentication (SharePoint Server 2010) ........ 170

Plan authentication methods (SharePoint Server 2010) ........ 171
- Supported authentication methods ........ 171
- Authentication modes — classic or claims-based ........ 172
- Implementing Windows authentication ........ 176
- Implementing forms-based authentication ........ 177
- Implementing SAML token-based authentication ........ 178
- Choosing authentication for LDAP environments ........ 180
- Planning zones for Web applications ........ 180

Architecture for SAML token-based providers .......... 183

Plan the Secure Store Service (SharePoint Server 2010) .......... 187
About the Secure Store Service .......... 187
Secure store service preparation .......... 187
Application IDs .......... 188
Secure store service mappings .......... 188
Secure store service and claims authentication .......... 188

Plan security hardening (SharePoint Server 2010) .......... 189
Secure server snapshots .......... 189
Web server and application server roles .......... 189
Database server role .......... 192
Specific port, protocol, and service guidance .......... 193
Blocking the standard SQL Server ports .......... 193
Configuring SQL Server database instances to listen on a nonstandard port .......... 194
Configuring SQL Server client aliases .......... 194
Service application communication .......... 194
File and Printer Sharing service requirements .......... 195
User Profile service hardening requirements .......... 196
Connections to external servers .......... 196
Service requirements for e-mail integration .......... 198
SMTP service .......... 198
Microsoft SharePoint Directory Management service .......... 198
Service requirements for session state .......... 198
SharePoint 2010 Products services .......... 198
Web.config file .......... 199

Plan automatic password change (SharePoint Server 2010) .......... 200
Configuring managed accounts .......... 200
Resetting passwords automatically on a schedule .......... 200
Detecting password expiration .......... 201
Resetting the account password immediately .......... 201
Synchronizing SharePoint Foundation account passwords with Active Directory Domain Services. 201
Resetting all passwords immediately .......... 201
Credential change process .......... 201

SQL Server and storage (SharePoint Server 2010) .......... 203

Overview of SQL Server in a SharePoint environment (SharePoint Server 2010) .......... 204
SharePoint 2010 Products and the SQL Server database engine .......... 204
Working with the SQL Server databases that support SharePoint 2010 Products .......... 205
SQL Server as a data platform for business intelligence in SharePoint 2010 Products .......... 205
SQL Server database engine .......... 205

SQL Server Analysis Services (SSAS): multi-dimensional data ........ 205
SQL Server Analysis Services: data mining ........ 206
SQL Server Reporting Services (SSRS) ........ 206
SQL Server Integration Services (SSIS) ........ 206
Business Intelligence Development Studio (BIDS) ........ 207
PowerPivot for Excel and PowerPivot for SharePoint ........ 207
Master Data Services ........ 207
StreamInsight and complex event processing ........ 207
SharePoint Server 2010 authoring and publishing tools for business intelligence ........ 208
Related content ........ 208

SQL Server 2008 R2 and SharePoint 2010 Products: Better Together (white paper) (SharePoint Server 2010) ........ 210

Storage and SQL Server capacity planning and configuration (SharePoint Server 2010) ........ 211
Design and configuration process for SharePoint 2010 Products storage and database tier ........ 211
Gather storage and SQL Server space and I/O requirements ........ 211
Databases used by SharePoint 2010 Products ........ 212
Understand SQL Server and IOPS ........ 213
Estimate core storage and IOPS needs ........ 214
Configuration storage and IOPS ........ 214
Content storage and IOPS ........ 214
Estimate service application storage needs and IOPS ........ 216
SharePoint Foundation 2010 service application storage and IOPS requirements ........ 217
SharePoint Server 2010 service application storage and IOPs requirements ........ 218
Determine availability needs ........ 220
Choose SQL Server version and edition ........ 220
Design storage architecture based on capacity and I/O requirements ........ 222
Choose a storage architecture ........ 222
Direct Attached Storage (DAS) ........ 222
Storage Area Network (SAN) ........ 222
Network Attached Storage (NAS) ........ 223
Choose disk types ........ 223
Choose RAID types ........ 223
Estimate memory requirements ........ 223
Understand network topology requirements ........ 224
Configure SQL Server ........ 225
Estimate how many servers are required ........ 225
Configure storage and memory ........ 225
Follow vendor storage configuration recommendations ........ 226
Provide as many resources as possible ........ 226
Set SQL Server options ........ 226

Configure databases ... 226
Separate and prioritize your data among disks ... 227
Use multiple data files for content databases ... 227
Limit content database size to improve manageability ... 228
Proactively manage the growth of data and log files ... 228
Validate and monitor storage and SQL Server performance ... 229
SQL Server counters to monitor ... 230
Physical server counters to monitor ... 231
Disk counters to monitor ... 232
Other monitoring tools ... 234

Overview of Remote BLOB Storage (SharePoint Server 2010) ... 235
Introduction to RBS ... 235
Using RBS together with SharePoint 2010 Products ... 236

Plan for Remote BLOB Storage (RBS) (SharePoint Server 2010) ... 238
Review the environment ... 238
Content database sizes ... 239
Content type and usage ... 239
Evaluate provider options ... 239

Plan for business continuity management (SharePoint Server 2010) ... 241
Business continuity management capabilities ... 241
Service level agreements ... 242
Related content ... 243

Plan to protect content by using recycle bins and versioning (SharePoint Server 2010) ... 245
Protecting content by using recycle bins ... 245
First-stage Recycle Bin ... 246
Second stage (Site Collection) Recycle Bin ... 246
Protecting content by using versioning ... 247

Plan for backup and recovery (SharePoint Server 2010) ... 248
Define business requirements ... 248
Choose what to protect and recover in your environment ... 249
Choose what to recover from within SharePoint content databases ... 251
Protecting customizations ... 251
Protecting workflows ... 251
Protecting service applications ... 252
Protecting SQL Server Reporting Services databases ... 253
Choose tools ... 253
Test hardware ... 253
Determine strategies ... 254
Plan for enhanced backup and recovery performance ... 255

Follow recommendations for configuring SQL Server and storage ... 256
Minimize latency between SQL Server and the backup location ... 256
Avoid processing conflicts ... 256
Follow SQL Server backup and restore optimization recommendations ... 256
Ensure sufficient write performance on the backup drive ... 257
Related content ... 257

Backup and recovery overview (SharePoint Server 2010) ... 258
Backup and recovery scenarios ... 258
Backup architecture ... 258
Farm backup architecture ... 258
Search service application backup process ... 261
Configuration-only backup use and benefits ... 261
Considerations for using farm backups ... 262
Granular backup and export architecture ... 263
Recovery processes ... 265
Restoring from a farm backup ... 265
Restoring as new versus restoring as overwrite ... 266
Search service application recovery process ... 267
Restoring from a site collection backup ... 267
Recovering from an unattached content database ... 267
Related content ... 268

Plan for availability (SharePoint Server 2010) ... 269
Availability overview ... 269
Costs of availability ... 270
Determining availability requirements ... 270
Choosing an availability strategy and level ... 271
Hardware component fault tolerance ... 271
Redundancy within a farm ... 271
Database availability strategies ... 273
Service application redundancy strategies ... 275
Search redundancy strategies within a farm ... 277
Redundancy and failover between closely located data centers configured as a single farm ("stretched" farm) ... 280

Plan for disaster recovery (SharePoint Server 2010) ... 282
Disaster recovery overview ... 282
Choose a disaster recovery strategy ... 283
Planning for cold standby data centers ... 283
Planning for warm standby data centers ... 284
Planning for hot standby data centers ... 284
Service application redundancy across data centers ... 286

Databases that can be log-shipped or asynchronously mirrored ............................................... 286
Service applications and databases that cannot be log-shipped or asynchronously mirrored ... 288
System requirements for disaster recovery............................................................................ 289

Global deployment of multiple farms (SharePoint Server 2010) .......................................... 290

Global solutions for SharePoint 2010 Products (model) ....................................................... 291

Client solutions for WAN environments (SharePoint Server 2010) ....................................... 292
Mobile views ......................................................................................................................... 293
Office Web Apps ................................................................................................................... 293
Office 2010 Document Cache and the MS-FSSHTTP protocol .......................................... 295
Outlook 2010 ......................................................................................................................... 297
SharePoint Workspace......................................................................................................... 298
SharePoint Workspace Mobile for Windows Phone 7 ......................................................... 299
SharePoint Workspace with Groove Server.......................................................................... 300

Planning worksheets for SharePoint Server 2010 ................................................................... 301
Planning worksheets by task................................................................................................ 301
Planning worksheets by title ................................................................................................. 304

# Getting help

Every effort has been made to ensure the accuracy of this book. This content is also available online in the Office System TechNet Library, so if you run into problems you can check for updates at:

*http://technet.microsoft.com/office*

If you do not find your answer in our online content, you can send an e-mail message to the Microsoft Office System and Servers content team at:

*itspdocs@microsoft.com*

If your question is about Microsoft Office products, and not about the content of this book, please search the Microsoft Help and Support Center or the Microsoft Knowledge Base at:

*http://support.microsoft.com*

# Technical diagrams (SharePoint Server 2010)

Many of these resources are visual representations of recommended solutions. They include poster-sized documents available in formats including Microsoft Office Visio 2007 or Microsoft Visio 2010 files (.vsd), PDF files, and XPS files. You might need extra software to view these files. See the following table for information about opening these files.

| File type | Software |
|---|---|
| .vsd | Office Visio 2007, Microsoft Visio 2010, or the free Visio viewer (http://go.microsoft.com/fwlink/?LinkId=118761)<br>If you use the Visio viewer, right-click the VSD link, click **Save Target As**, save the file to your computer, and then open the file from your computer. |
| .pdf | Any PDF viewer, such as Adobe Reader (http://go.microsoft.com/fwlink/?LinkId=134751) |
| .xps | Windows 7, Windows Vista, Windows XP with .NET Framework 3.0, or XPS Essentials Pack (http://go.microsoft.com/fwlink/?LinkId=134750) |

## Models

Models are 34-by-44-inch posters that detail a specific technical area. These models are intended to be used with corresponding articles on TechNet. These models are created by using Office Visio 2007. You can modify the Visio files to illustrate how you plan to incorporate Microsoft SharePoint 2010 Products in your own environment.

| Title | Description |
|---|---|
| **Design Sample: Corporate Portal with Classic Authentication** | Illustrate a typical corporate deployment, with the most common types of sites represented. The two samples differ only in the mode of authentication that is implemented.<br>Use these design samples with the |

| Title | Description |
|---|---|
| 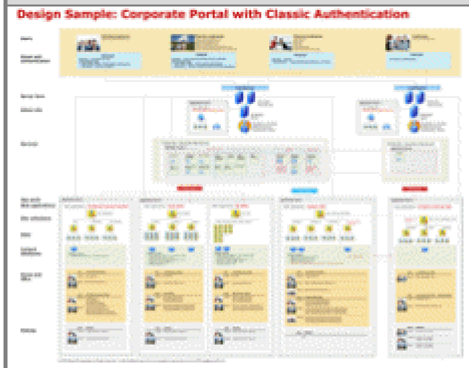 <br>Visio (http://go.microsoft.com/fwlink/?LinkId=196969)<br>PDF (http://go.microsoft.com/fwlink/?LinkId=196970)<br>XPS (http://go.microsoft.com/fwlink/?LinkId=196971)<br>**Design Sample: Corporate Portal with Claims-based Authentication**<br> <br>Visio (http://go.microsoft.com/fwlink/?LinkId=196972)<br>PDF (http://go.microsoft.com/fwlink/?LinkId=196973)<br>XPS (http://go.microsoft.com/fwlink/?LinkId=196974) | following article: Design sample: Corporate deployment (SharePoint Server 2010) |
| **SharePoint 2010 Products Deployment** | Presents such deployment-related information as the different deployment stages and environments, plus a flowchart that illustrates the steps for installing and configuring SharePoint 2010 Products. |

| Title | Description |
|---|---|
| 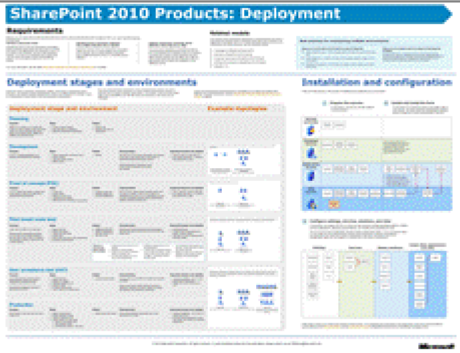<br>SharePoint 2010 Products: Deployment<br>Requirements<br>Deployment stages and environments<br>Installation and configuration<br>Visio (http://go.microsoft.com/fwlink/?LinkId=183024)<br>PDF (http://go.microsoft.com/fwlink/?LinkId=183025)<br>XPS (http://go.microsoft.com/fwlink/?LinkId=183026) | |
| **Services in SharePoint 2010 Products**<br>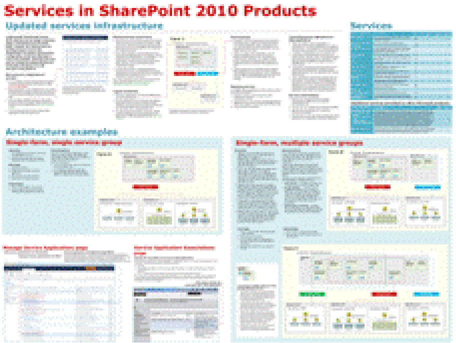<br>Services in SharePoint 2010 Products<br>Updated services infrastructure<br>Services<br>Architecture examples<br>Visio (http://go.microsoft.com/fwlink/?LinkID=167090)<br>PDF (http://go.microsoft.com/fwlink/?LinkID=167092)<br>XPS (http://go.microsoft.com/fwlink/?LinkID=167091) | Describes and illustrates the services architecture, including common ways to deploy services in your overall solution design.<br>Use this diagram with the following articles:<br>• Services architecture planning (SharePoint Foundation 2010)<br>• Services architecture planning (SharePoint Server 2010) |
| **Cross-farm Services in SharePoint 2010 Products**<br>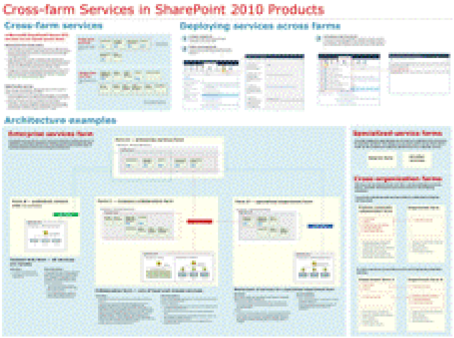<br>Cross-farm Services in SharePoint 2010 Products<br>Cross-farm services<br>Deploying services across farms<br>Architecture examples | Illustrates how to deploy services across farms to provide centralized administration of services.<br>Use this diagram with the following articles:<br>• Services architecture planning (SharePoint Foundation 2010)<br>• Services architecture planning (SharePoint Server 2010) |

| Title | Description |
|---|---|
| Visio (http://go.microsoft.com/fwlink/?LinkID=167093)<br>PDF (http://go.microsoft.com/fwlink/?LinkID=167095)<br>XPS (http://go.microsoft.com/fwlink/?LinkID=167094) | |
| **Topologies for SharePoint Server 2010**<br>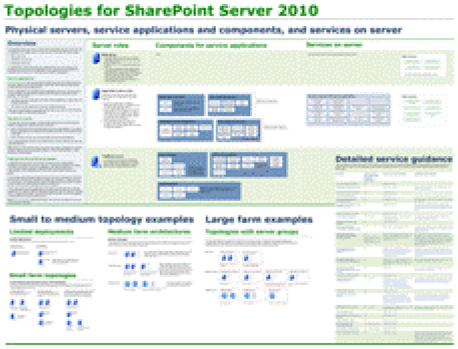<br> Topologies for SharePoint Server 2010<br>Physical servers, service applications and components, and services on server<br>Detailed service guidance<br>Small to medium topology examples<br>Large farm examples<br>Visio (http://go.microsoft.com/fwlink/?LinkID=167087)<br>PDF (http://go.microsoft.com/fwlink/?LinkID=167089)<br>XPS (http://go.microsoft.com/fwlink/?LinkID=167088) | Describes common ways to build and scale farm topologies, including planning which servers to start services on. |
| **Extranet Topologies for SharePoint 2010 Products**<br>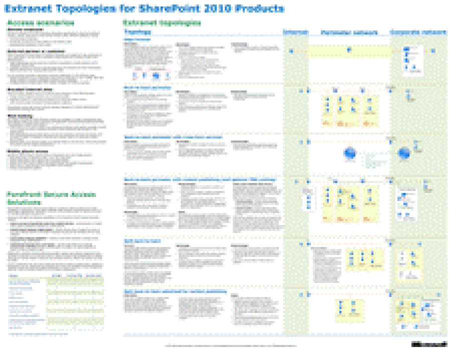<br> Extranet Topologies for SharePoint 2010 Products<br>Access scenarios<br>Extranet topologies<br>Forefront Secure Access Solutions<br>Visio (http://go.microsoft.com/fwlink/?LinkId=187987)<br>PDF (http://go.microsoft.com/fwlink/?LinkId=187988)<br>XPS (http://go.microsoft.com/fwlink/?LinkId=187986) | Illustrates the specific extranet topologies that have been tested with SharePoint 2010 Products. Provides a comparison of ISA Server, Forefront TMG, Forefront UAG when used as a firewall or gateway product with SharePoint 2010 Products. |
| **Hosting Environments in SharePoint 2010 Products** | Summarizes the support for hosting environments and illustrates common hosting architectures.<br>For more information on designing and |

| Title | Description |
|---|---|
| 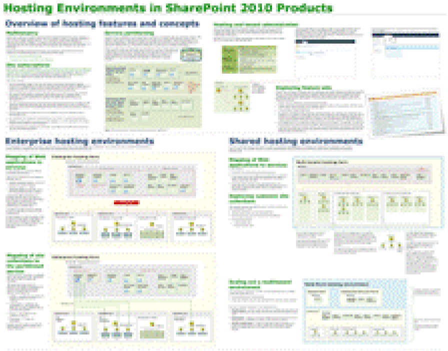  Hosting Environments in SharePoint 2010 Products<br>Overview of hosting features and concepts<br>Enterprise hosting environments<br>Shared hosting environments<br>Visio (http://go.microsoft.com/fwlink/?LinkID=167084)<br>PDF (http://go.microsoft.com/fwlink/?LinkID=167086)<br>XPS (http://go.microsoft.com/fwlink/?LinkID=167085) | deploying hosting environments, see the following: White paper: SharePoint 2010 for hosters (SharePoint Server 2010). |
| **Search Technologies for SharePoint 2010 Products**<br>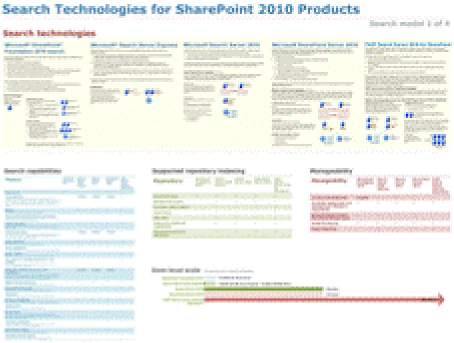  Search Technologies for SharePoint 2010 Products<br>Search technologies<br>Visio (http://go.microsoft.com/fwlink/?LinkID=167731)<br>PDF (http://go.microsoft.com/fwlink/?LinkID=167733)<br>XPS (http://go.microsoft.com/fwlink/?LinkID=167732) | Compares and contrasts the search technologies that work with SharePoint Products 2010:<br>• SharePoint Foundation 2010<br>• Search Server 2010 Express<br>• Search Server 2010<br>• SharePoint Server 2010<br>• FAST Search Server 2010 for SharePoint |
| **Search Environment Planning for Microsoft SharePoint Server 2010** | Walks through primary architecture design decisions for search environments. |

| Title | Description |
|---|---|
| Search Environment Planning for Microsoft SharePoint Server 2010<br>Search model 2 of 4<br>1 Where are your users and content?<br>2 What content do you plan to include in search results?<br>3 How many farms do you have to plan for?<br>Visio (http://go.microsoft.com/fwlink/?LinkID=167734)<br>PDF (http://go.microsoft.com/fwlink/?LinkID=167736)<br>XPS (http://go.microsoft.com/fwlink/?LinkID=167735) | |
| **Search Architectures for Microsoft SharePoint Server 2010**<br>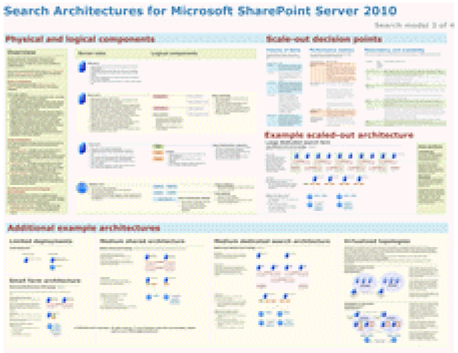 Search Architectures for Microsoft SharePoint Server 2010<br>Search model 3 of 4<br>Physical and logical components<br>Scale-out decision points<br>Example scaled-out architecture<br>Additional example architectures<br>Visio (http://go.microsoft.com/fwlink/?LinkID=167737)<br>PDF (http://go.microsoft.com/fwlink/?LinkID=167739)<br>XPS (http://go.microsoft.com/fwlink/?LinkID=167738) | Details the physical and logical architecture components that make up a search system and illustrates common search architectures. |
| **Design Search Architectures for Microsoft SharePoint Server 2010** | Walks through the initial design steps to determine a basic design for a SharePoint Server 2010 search architecture. |

| Title | Description |
|---|---|
| 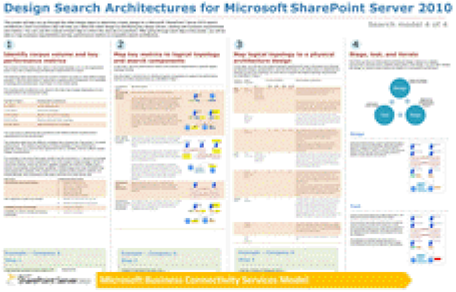 Design Search Architectures for Microsoft SharePoint Server 2010<br>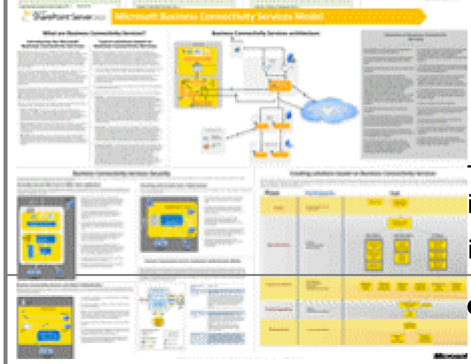 <br>.inkID=167740)<br>inkID=167742)<br>inkID=167741) | |
| **del**<br>Visio (http://go.microsoft.com/fwlink/?LinkId=165565)<br>PDF (http://go.microsoft.com/fwlink/?LinkID=165566)<br>XPS (http://go.microsoft.com/fwlink/?LinkId=165571) | Microsoft Business Connectivity Services are a set of services and features in Microsoft SharePoint Server 2010 and Microsoft SharePoint Foundation 2010 that support integrating data from external systems into solutions based on Microsoft SharePoint Server and Microsoft SharePoint Foundation. This model poster describes the architecture of Microsoft Business Connectivity Services in SharePoint Server 2010 and provides information about how to create solutions that are based on the service.<br>Use this model with the following article: Business Connectivity Services overview (SharePoint Server 2010) |
| **Content Deployment in SharePoint Server 2010** | Describes the content deployment feature in SharePoint Server 2010. It includes information about the following:<br>• Overview of content deployment<br>• Description of content deployment paths and jobs |

| Title | Description |
|---|---|
| 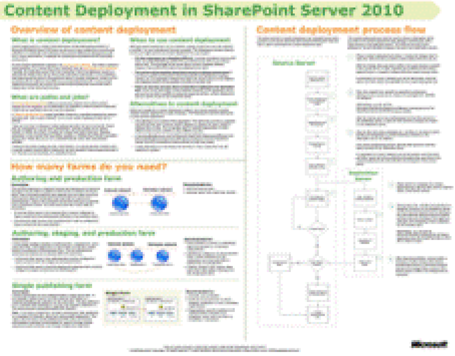<br>Content Deployment in SharePoint Server 2010<br>[Visio](http://go.microsoft.com/fwlink/?LinkID=179391&clcid=0x409) (http://go.microsoft.com/fwlink/?LinkID=179391&clcid=0x409)<br>[PDF](http://go.microsoft.com/fwlink/?LinkID=179523&clcid=0x409) (http://go.microsoft.com/fwlink/?LinkID=179523&clcid=0x409)<br>[XPS](http://go.microsoft.com/fwlink/?LinkID=179524&clcid=0x409) (http://go.microsoft.com/fwlink/?LinkID=179524&clcid=0x409) | • When to use content deployment<br>• Alternatives to content deployment<br>• Illustrates common content deployment farm topologies<br>• Illustrates and explains the overall content deployment process |
| **Microsoft SharePoint Server 2010 Upgrade Planning**<br>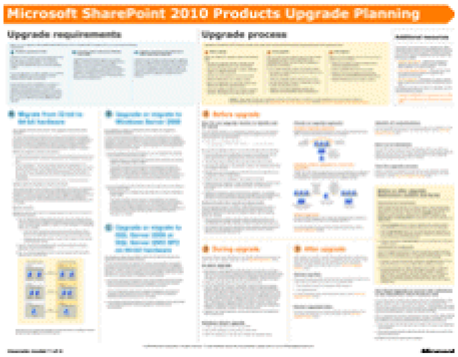<br>Microsoft SharePoint 2010 Products Upgrade Planning<br>[Visio](http://go.microsoft.com/fwlink/?LinkId=167098) (http://go.microsoft.com/fwlink/?LinkId=167098)<br>[PDF](http://go.microsoft.com/fwlink/?LinkId=167099) (http://go.microsoft.com/fwlink/?LinkId=167099)<br>[XPS](http://go.microsoft.com/fwlink/?LinkId=167100) (http://go.microsoft.com/fwlink/?LinkId=167100) | Covers planning for an upgrade from Microsoft Office SharePoint Server 2007 to SharePoint Server 2010. It includes information about the following:<br>• Upgrade requirements: Hardware, operating system, and database<br>• Upgrade process: specific steps to follow before, during, and after the upgrade<br>Use this model with the following article: [Upgrading to SharePoint Server 2010]() |
| **Microsoft SharePoint Server 2010 Upgrade Approaches** | Helps you understand the in-place, database attach, and hybrid approaches to upgrading from Office SharePoint Server 2007 to SharePoint Server 2010.<br>• See the farm topologies before, during, and after upgrade |

| Title | Description |
|---|---|
| 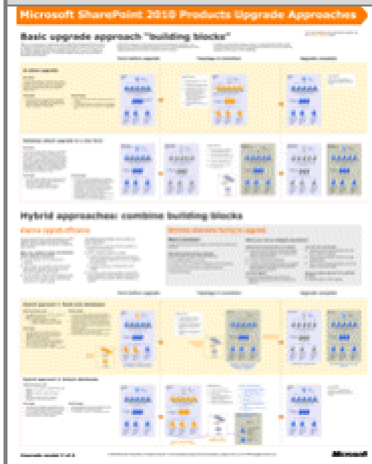<br>Microsoft SharePoint 2010 Products Upgrade Approaches<br>Basic upgrade approach "building blocks"<br>Hybrid approaches: combine building blocks<br>[Visio](http://go.microsoft.com/fwlink/?LinkId=167101) (http://go.microsoft.com/fwlink/?LinkId=167101)<br>[PDF](http://go.microsoft.com/fwlink/?LinkId=167102) (http://go.microsoft.com/fwlink/?LinkId=167102)<br>[XPS](http://go.microsoft.com/fwlink/?LinkId=167103) (http://go.microsoft.com/fwlink/?LinkId=167103) | • Compare the advantages of each type of upgrade approach<br>Use this model with the following articles:<br>• [Determine upgrade approach (SharePoint Server 2010)]()<br>• [Upgrade process overview (SharePoint Server 2010)]() |
| **Microsoft SharePoint Server 2010 — Test Your Upgrade Process**<br>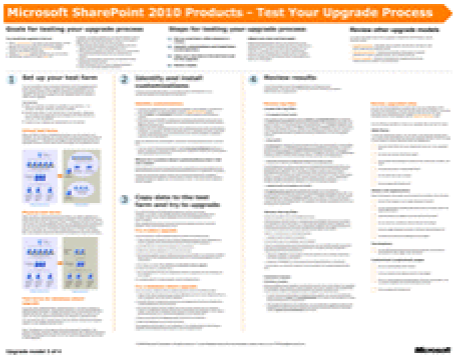<br>Microsoft SharePoint 2010 Products - Test Your Upgrade Process<br>[Visio](http://go.microsoft.com/fwlink/?LinkId=167104) (http://go.microsoft.com/fwlink/?LinkId=167104)<br>[PDF](http://go.microsoft.com/fwlink/?LinkId=167105) (http://go.microsoft.com/fwlink/?LinkId=167105)<br>[XPS](http://go.microsoft.com/fwlink/?LinkId=167106) (http://go.microsoft.com/fwlink/?LinkId=167106) | Explains the methodology for testing the upgrade process before upgrading from Office SharePoint Server 2007 to SharePoint Server 2010.<br>• Understand the goals for testing your upgrade process: customizations, hardware, timing, planning<br>• See specific steps to follow for testing your upgrade process<br>Use this model with the following article: [Use a trial upgrade to find potential issues (SharePoint Server 2010)]() |
| **Microsoft SharePoint Server 2010 — Services Upgrade** | Covers upgrading services from Office SharePoint Server 2007 to SharePoint Server 2010.<br>• Considerations for specific services: Personalization, Search, InfoPath Forms, Excel, Business Data |

| Title | Description |
|---|---|
| Microsoft SharePoint Server 2010 Services Upgrade<br>Before you begin<br>Considerations for specific services<br>In-place upgrade with services<br>Database attach upgrade with services<br>[Visio](http://go.microsoft.com/fwlink/?LinkId=167107) (http://go.microsoft.com/fwlink/?LinkId=167107)<br>[PDF](http://go.microsoft.com/fwlink/?LinkId=167108) (http://go.microsoft.com/fwlink/?LinkId=167108)<br>[XPS](http://go.microsoft.com/fwlink/?LinkId=167109) (http://go.microsoft.com/fwlink/?LinkId=167109) | Catalog, Single Sign-on<br>• In-place upgrade with services<br>• Database attach upgrade with services |
| **Microsoft SharePoint Server 2010 — Upgrading Parent and Child Farms**<br>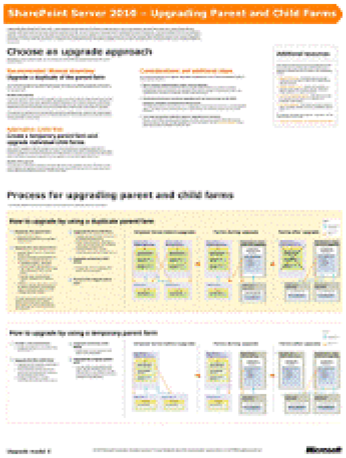 SharePoint Server 2010 – Upgrading Parent and Child Farms<br>Choose an upgrade approach<br>Process for upgrading parent and child farms<br>[Visio](http://go.microsoft.com/fwlink/?LinkId=190984) (http://go.microsoft.com/fwlink/?LinkId=190984)<br>[PDF](http://go.microsoft.com/fwlink/?LinkId=190985) (http://go.microsoft.com/fwlink/?LinkId=190985)<br>[XPS](http://go.microsoft.com/fwlink/?LinkId=190986) (http://go.microsoft.com/fwlink/?LinkId=190986) | Covers the process for and considerations to keep in mind when you upgrade farms that share services (parent and child farms). |
| **Getting started with business intelligence in SharePoint Server 2010** | Covers an overview of business intelligence in SharePoint Server and provides you with the following |

| Title | Description |
|---|---|
| 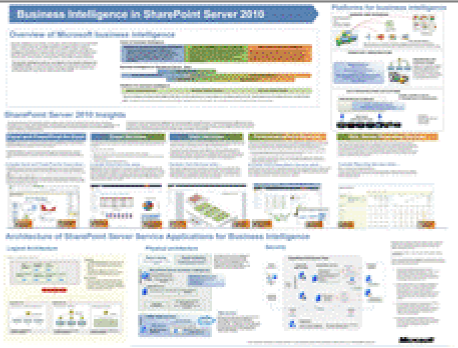 <br>Visio (http://go.microsoft.com/fwlink/?LinkId=167409)<br>PDF (http://go.microsoft.com/fwlink/?LinkId=167170)<br>XPS (http://go.microsoft.com/fwlink/?LinkId=167171) | information.<br>• An overview of each business intelligence service and when you might use the service.<br>• Architecture for application of the business intelligence services and how they work together in a topology.<br>• A list of possible data sources for each business intelligence service. |
| **Databases That Support SharePoint 2010 Products**<br>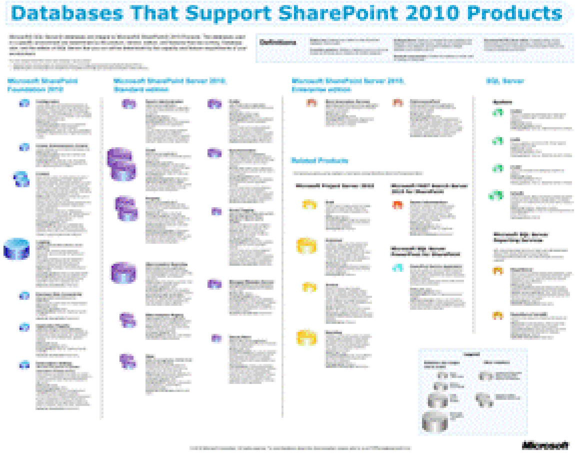 <br>Visio (http://go.microsoft.com/fwlink/?LinkId=187970)<br>PDF (http://go.microsoft.com/fwlink/?LinkId=187969)<br>XPS (http://go.microsoft.com/fwlink/?LinkId=187971) | Describes the Microsoft SQL Server databases on which SharePoint Server 2010 runs. |
| **SharePoint 2010 Products: Virtualization Process** | Provides guidance related to virtualization and the various stages of deployment, as well as requirements and examples.<br>Use this diagram with the articles in the following chapters:<br>• Virtualization planning (SharePoint |

| Title | Description |
|---|---|
| <br>SharePoint 2010 Products: Virtualization Process<br>Visio (http://go.microsoft.com/fwlink/?LinkId=195021)<br>PDF (http://go.microsoft.com/fwlink/?LinkId=195022)<br>XPS (http://go.microsoft.com/fwlink/?LinkId=195023) | Foundation 2010)<br>- Virtualization planning (SharePoint Server 2010) |
| **Governance for SharePoint Server 2010**<br>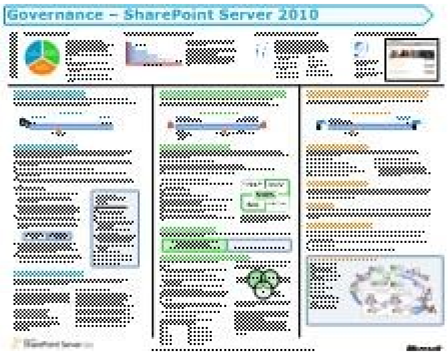<br>Governance – SharePoint Server 2010<br>Visio (http://go.microsoft.com/fwlink/?LinkId=200532)<br>PDF (http://go.microsoft.com/fwlink/?LinkId=200533)<br>XPS (http://go.microsoft.com/fwlink/?LinkId=200534) | Illustrates how to develop a governance plan that includes IT governance, information management governance, and application management governance.<br>Use this diagram with the following articles:<br>- Governance overview (SharePoint Server 2010)<br>- Governance features (SharePoint Server 2010) |
| **Duet Enterprise for Microsoft SharePoint and SAP Poster** | Illustrates Duet Enterprise architecture for both the SAP and Microsoft environments, with detailed explanations of each area. |

| Title | Description |
|---|---|
| 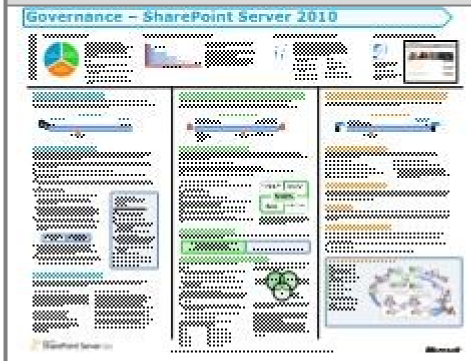<br>Governance – SharePoint Server 2010<br>Visio<br>(http://go.microsoft.com/fwlink/?LinkID=208107&clcid=0x409)<br>PDF<br>(http://go.microsoft.com/fwlink/?LinkID=208108&clcid=0x409)<br>XPS<br>(http://go.microsoft.com/fwlink/?LinkId=208109&clcid=0x409) | |

# Tips for printing posters

If you have a plotter, you can print these posters in their full size. If you don't have plotter, use the following steps to print on smaller paper.

### Print posters on smaller paper

1. Open the poster in Visio.
2. On the **File** menu, click **Page Setup**.
3. On the **Print Setup** tab, in the **Printer paper** section, select the size of paper you want to print on.
4. On the **Print Setup** tab, in the **Print zoom** section, click **Fit to**, and then enter **1 sheet across by 1 sheet down**.
5. On the **Page Size** tab, click **Size to fit drawing contents**, and then click **OK**.
6. On the **File** menu, click **Print**.

If you want to create posters that use the same symbols as these posters, you can download Visio stencils for posters (http://www.microsoft.com/downloads/en/details.aspx?FamilyID=88e03d22-8f42-4c9d-94ef-d8e48322d677).

# Plan for server farms and environments (SharePoint Server 2010)

This section contains infrastructure planning resources and articles designed to help you plan server farms and environments required to support your SharePoint sites.

In this section:

- System requirements (SharePoint Server 2010)
- Services architecture planning (SharePoint Server 2010)
- Logical architecture components (SharePoint Server 2010)
- Plan authentication (SharePoint Server 2010)
- Plan security hardening (SharePoint Server 2010)
- Plan for business continuity management (SharePoint Server 2010)
- Performance and capacity management (SharePoint Server 2010)
- Virtualization planning (SharePoint Server 2010)

# System requirements (SharePoint Server 2010)

Before you install Microsoft SharePoint Server 2010, you must ensure that you have installed all required hardware and software. To effectively plan your deployment, you must understand the level of support provided for the Web browsers that you will be using in your environment and how support for IP versions 4 and 6 is implemented in SharePoint Server 2010. You must also understand the URL and path length restrictions in SharePoint Server 2010.

The articles in this section help you prepare for the installation of SharePoint Server 2010 by providing information about the prerequisites that you need to run SharePoint Server 2010.

- Hardware and software requirements (SharePoint Server 2010)

  This article describes the hardware and software requirements that you must meet to successfully install SharePoint Server 2010.

- Plan browser support (SharePoint Server 2010)

  This article describes levels of support for Web browsers to use with SharePoint Server 2010.

- URL path length restrictions (SharePoint Server 2010)

  This article discusses the specific URL path length and character restrictions in SharePoint Server 2010, Internet Explorer 7, and Internet Explorer 8 that you should be aware of when planning sites, navigation, and structure.

- IP support (SharePoint Server 2010)

  This article describes SharePoint Server 2010 support for IP version 4 (IPv4) and IP version 6 (IPv6).

- Windows Server 2008 R2 and SharePoint Server 2010: Better Together (white paper)

  This article describes the benefits of deploying SharePoint Server 2010 on Windows Server 2008 R2 Enterprise.

- SQL Server 2008 R2 and SharePoint 2010 Products: Better Together (white paper) (SharePoint Server 2010)

  This article describes the benefits of deploying SharePoint Server 2010 on Microsoft SQL Server 2008 R2 Enterprise, and provides a comparison of the functionality available in SharePoint when it is running on different versions and editions of SQL Server.

- Business Productivity at Its Best: Microsoft Office 2010 and SharePoint Server 2010 Better Together (white paper)

  This article describes the benefits of using Microsoft Office 2010 with SharePoint Server 2010.

# Hardware and software requirements (SharePoint Server 2010)

This article lists the minimum hardware and software requirements to install and run Microsoft SharePoint Server 2010.

**Important:**
If you contact Microsoft technical support about a production system that does not meet the minimum hardware specifications described in this document, support will be limited until the system is upgraded to the minimum requirements.

In this article:

- Overview
- Hardware requirements—Web servers, application servers, and single server installations
- Hardware requirements—Database servers
- Software requirements
- Access to applicable software

## Overview

Microsoft SharePoint Server 2010 provides for a number of installation scenarios. Currently, these installations include single server with built-in database installations and single-server or multiple-server farm installations.

If you plan on installing Microsoft Project Server 2010 with SharePoint Server 2010, see Hardware and software requirements (Project Server 2010). Especially note the supported Web browsers for Project Web App users.

## Hardware requirements—Web servers, application servers, and single server installations

The requirements in the following table apply both to installations on a single server with a built-in database and to servers running SharePoint Server 2010 in a multiple server farm installation.

| **Component** | **Minimum requirement** |
|---|---|
| Processor | 64-bit, four cores |
| RAM | • 4 GB for developer or evaluation use<br>• 8 GB for production use in a single server or |

| Component | Minimum requirement |
|---|---|
| | multiple server farm |
| Hard disk | 80 GB for system drive<br>For production use, you need additional free disk space for day-to-day operations. Maintain twice as much free space as you have RAM for production environments. For more information, see Capacity management and sizing for SharePoint Server 2010. |

## Hardware requirements—Database servers

The requirements in the following table apply to database servers in production environments with multiple servers in the farm.

**Note:**
Our definitions of small and medium deployments are those described in the "Reference Architectures" section in Capacity management and sizing for SharePoint Server 2010.

| Component | Minimum requirement |
|---|---|
| Processor | • 64-bit, four cores for small deployments<br>• 64-bit, eight cores for medium deployments |
| RAM | • 8 GB for small deployments<br>• 16 GB for medium deployments<br>For large deployments, see the "Estimate memory requirements" section in Storage and SQL Server capacity planning and configuration (SharePoint Server 2010).<br>**Note:**<br>These values are higher than those recommended as the minimum values for SQL Server because of the distribution of data required for a SharePoint Products 2010 environment. For more information about SQL Server system requirements, see Hardware and Software Requirements for Installing SQL |

| Component | Minimum requirement |
|---|---|
| | Server 2008 (http://go.microsoft.com/fwlink/?LinkId=129377). |
| Hard disk | 80 GB for system drive<br>Hard disk space is dependent on the size of your SharePoint content. For information about estimating the size of content and other databases for your deployment, see Storage and SQL Server capacity planning and configuration (SharePoint Server 2010). |

# Software requirements

The requirements in the following tables apply to single server with built-in database installations and server farm installations that include a single server and multiple servers in the farm.

**Important:**
SharePoint Server 2010 does not support single label domain names. For more information, see Information about configuring Windows for domains with single-label DNS names.

The Microsoft SharePoint Products Preparation Tool — which you access from the SharePoint Server 2010 Start page — can assist you in the installation of the software prerequisites for SharePoint Server 2010. Ensure that you have an Internet connection, because some of these prerequisites are installed from the Internet. For more information, see Deploy a single server with SQL Server (SharePoint Server 2010), Deploy a single server with a built-in database (SharePoint Server 2010), and Multiple servers for a three-tier farm (SharePoint Server 2010).

## Minimum requirements

| Environment | Minimum requirement |
|---|---|
| Database server in a farm | One of the following:<br>• The 64-bit edition of Microsoft SQL Server 2008 R2.<br>• The 64-bit edition of Microsoft SQL Server 2008 with Service Pack 1 (SP1) and Cumulative Update 2. From the Cumulative update package 2 for SQL Server 2008 Service Pack 1 (http://go.microsoft.com/fwlink/?LinkId=165962) page, click the **View and request hotfix downloads** link and |

| Environment | Minimum requirement |
|---|---|
| | follow the instructions. On the Hotfix Request page, download the SQL_Server_2008_SP1_Cumulative_Update_2 file. When you install Microsoft SQL Server 2008 SP1 on Windows Server 2008 R2, you might receive a compatibility warning. You can disregard this warning and continue with your installation.<br>**Note:**<br>We do not recommend that you use CU3 or CU4, but instead CU2, CU5, or a later CU than CU5. For more information, see Cumulative update package 5 for SQL Server 2008 (http://go.microsoft.com/fwlink/?LinkId=196928). Download the SQL_Server_2008_RTM_CU5_SNAC file.<br>• The 64-bit edition of Microsoft SQL Server 2005 with Service Pack 3 (SP3). From the Cumulative update package 3 for SQL Server 2005 Service Pack 3 (http://go.microsoft.com/fwlink/?LinkId=165748) page, click the **View and request hotfix downloads** link and follow the instructions. On the Hotfix Request page, download the SQL_Server_2005_SP3_Cumulative_Update_3 file.<br>For more information about choosing a version of SQL Server, see SQL Server 2008 R2 and SharePoint 2010 Products: Better Together (white paper) (SharePoint Server 2010). |
| Single server with built-in database | • The 64-bit edition of Windows Server 2008 Standard, Enterprise, Data Center, or Web Server with SP2, or the 64-bit edition of Windows Server 2008 R2 Standard, Enterprise, Data Center, or Web Server. If you are running Windows Server 2008 without SP2, the Microsoft SharePoint Products Preparation Tool installs Windows Server 2008 SP2 automatically.<br>**Note:**<br>You must download an update for Windows Server 2008 and Windows Server 2008 R2 before you run Setup. The update is a hotfix for |

| Environment | Minimum requirement |
|---|---|
| | the .NET Framework 3.5 SP1 that is installed by the Preparation tool. It provides a method to support token authentication without transport security or message encryption in WCF. For more information and links, see the "Access to Applicable Software" section later in this article.<br>• KB979917 - QFE for Sharepoint issues - Perf Counter fix & User Impersonation (http://go.microsoft.com/fwlink/?LinkId=192577)<br>• For Windows Server 2008 SP2, download the Windows6.0-KB979917-x64.msu (Vista) file.<br>• For Windows Server 2008 R2, download the Windows6.1-KB979917-x64.msu (Win7) file.<br>For information, see the related KB article Two issues occur when you deploy an ASP.NET 2.0-based application on a server that is running IIS 7.0 or IIS 7.5 in Integrated mode (http://go.microsoft.com/fwlink/?LinkId=192578).<br>**The preparation tool installs the following prerequisites:**<br>• Web Server (IIS) role<br>• Application Server role<br>• Microsoft .NET Framework version 3.5 SP1<br>• SQL Server 2008 Express with SP1<br>• Microsoft Sync Framework Runtime v1.0 (x64)<br>• Microsoft Filter Pack 2.0<br>• Microsoft Chart Controls for the Microsoft .NET Framework 3.5<br>• Windows PowerShell 2.0<br>• SQL Server 2008 Native Client<br>• Microsoft SQL Server 2008 Analysis Services ADOMD.NET<br>• ADO.NET Data Services Update for .NET Framework 3.5 SP1<br>• A hotfix for the .NET Framework 3.5 SP1 that provides a method to support token authentication without transport security or message encryption in WCF. |

| Environment | Minimum requirement |
|---|---|
|  | • Windows Identity Foundation (WIF)<br>**Note:**<br>If you have Microsoft "Geneva" Framework installed, you must uninstall it before you install the Windows Identity Foundation (WIF). |
| Front-end Web servers and application servers in a farm | • The 64-bit edition of Windows Server 2008 Standard, Enterprise, Data Center, or Web Server with SP2, or the 64-bit edition of Windows Server 2008 R2 Standard, Enterprise, Data Center, or Web Server. If you are running Windows Server 2008 with SP1, the Microsoft SharePoint Products Preparation Tool installs Windows Server 2008 SP2 automatically.<br>**Note:**<br>You must download an update for Windows Server 2008 and Windows Server 2008 R2 before you run Setup. The update is a hotfix for the .NET Framework 3.5 SP1 that is installed by the Preparation tool. It provides a method to support token authentication without transport security or message encryption in WCF. For more information and links, see the "Access to Applicable Software" section.<br>• KB979917 - QFE for Sharepoint issues - Perf Counter fix & User Impersonation (http://go.microsoft.com/fwlink/?LinkId=192577)<br>  • For Windows Server 2008 SP2, download the Windows6.0-KB979917-x64.msu (Vista) file.<br>  • For Windows Server 2008 R2, download the Windows6.1-KB979917-x64.msu (Win7) file.<br>For information, see the related KB article Two issues occur when you deploy an ASP.NET 2.0-based application on a server that is running IIS 7.0 or IIS 7.5 in Integrated mode (http://go.microsoft.com/fwlink/?LinkId=192578).<br>**The preparation tool installs the following prerequisites:**<br>• Web Server (IIS) role |

| Environment | Minimum requirement |
|---|---|
|  | - Application Server role<br>- Microsoft .NET Framework version 3.5 SP1<br>- Microsoft Sync Framework Runtime v1.0 (x64)<br>- Microsoft Filter Pack 2.0<br>- Microsoft Chart Controls for the Microsoft .NET Framework 3.5<br>- Windows PowerShell 2.0<br>- SQL Server 2008 Native Client<br>- Microsoft SQL Server 2008 Analysis Services ADOMD.NET<br>- ADO.NET Data Services Update for .NET Framework 3.5 SP1<br>- A hotfix for the .NET Framework 3.5 SP1 that provides a method to support token authentication without transport security or message encryption in WCF.<br>- Windows Identity Foundation (WIF)<br>**Note:**<br>If you have Microsoft "Geneva" Framework installed, you must uninstall it before you install the Windows Identity Foundation (WIF). |
| Client computer | - A supported browser. For more information, see Plan browser support (SharePoint Server 2010). |

## Optional software

| Environment | Optional software |
|---|---|
| Single server with built-in database and front-end Web servers and application servers in a farm | - Microsoft SQL Server 2008 R2 to work with PowerPivot workbooks. For more information, see Microsoft SQL Server 2008 R2 (http://go.microsoft.com/fwlink/?LinkID=179611).<br>- Windows 7 or Windows Vista. For more |

| Environment | Optional software |
|---|---|
| | information, see Setting Up the Development Environment for SharePoint Server (http://go.microsoft.com/fwlink/?LinkID=164557).<br>- SQL Server Remote BLOB Store installation package from the Feature Pack for Microsoft SQL Server 2008 R2. For the download, go to the Download Center (http://go.microsoft.com/fwlink/?LinkID=177388).<br>**The preparation tool installs the following optional software:**<br>- Microsoft SQL Server 2008 R2 Reporting Services Add-in for Microsoft SharePoint Technologies 2010 (SSRS) to use Access Services for SharePoint Server 2010. For the download, go to the Download Center (http://go.microsoft.com/fwlink/?LinkID=192588).<br>- Microsoft Server Speech Platform to make phonetic name matching work correctly for SharePoint Search 2010. |
| Client computer | - Microsoft Office 2010 client. For more information, see Microsoft Office 2010 (http://go.microsoft.com/fwlink/?LinkId=195843).<br>- Microsoft Silverlight 3. |

# Access to applicable software

To install Windows Server 2008, Microsoft SQL Server, or SharePoint Server 2010, you can go to the Web sites listed in this section. You can install most software prerequisites through the SharePoint Server 2010 Start page. The software prerequisites are also available from Web sites listed in this section. The Web Server (IIS) role and the Application Server role can be enabled manually in Server Manager.

In scenarios where installing prerequisites directly from the Internet is not possible or not feasible, you can install the prerequisites from a network share. For more information, see Install prerequisites from a network share (SharePoint Server 2010).

- SharePoint Server 2010 Standard Trial (http://go.microsoft.com/fwlink/?LinkId=197413)
- SharePoint Server 2010 Enterprise Trial (http://go.microsoft.com/fwlink/?LinkId=197414)
- 2010 Server Language Packs for SharePoint Server 2010, Project Server 2010, Search Server 2010, and Office Web Apps 2010 (http://go.microsoft.com/fwlink/?LinkId=197415)
- Windows Server 2008 R2 and SharePoint Server 2010: Better Together (white paper)
- Business Productivity at Its Best: Microsoft Office 2010 and SharePoint Server 2010 Better Together (white paper)
- Windows Server 2008 (http://go.microsoft.com/fwlink/?LinkId=197426)
- Windows Server 2008 R2 (http://go.microsoft.com/fwlink/?LinkId=197428)
- SQL Server 2008 R2 (http://go.microsoft.com/fwlink/?LinkId=197429)
- SQL Server 2008 (http://go.microsoft.com/fwlink/?LinkID=179611)
- SQL Server 2005 (http://go.microsoft.com/fwlink/?LinkId=197431)
- Microsoft SQL Server 2008 SP1 (http://go.microsoft.com/fwlink/?LinkId=166490)
- Cumulative update package 2 for SQL Server 2008 Service Pack 1 (http://go.microsoft.com/fwlink/?LinkId=165962)
- Cumulative update package 5 for SQL Server 2008 (http://go.microsoft.com/fwlink/?LinkId=197434). Download the SQL_Server_2008_RTM_CU5_SNAC file.
- Microsoft SQL Server 2005 SP3 (http://go.microsoft.com/fwlink/?LinkId=166496)
- Cumulative update package 3 for SQL Server 2005 Service Pack 3 (http://go.microsoft.com/fwlink/?LinkId=165748)
- Microsoft Windows Server 2008 SP2 (http://go.microsoft.com/fwlink/?LinkId=166500)
- Windows Server 2008 with SP 2 FIX: A hotfix that provides a method to support the token authentication without transport security or message encryption in WCF is available for the .NET Framework 3.5 SP1 (http://go.microsoft.com/fwlink/?LinkID=160770)
- Windows Server 2008 R2 FIX: A hotfix that provides a method to support the token authentication without transport security or message encryption in WCF is available for the .NET Framework 3.5 SP1 (http://go.microsoft.com/fwlink/?LinkID=166231)
- Microsoft .NET Framework 3.5 Service Pack 1 (http://go.microsoft.com/fwlink/?LinkId=131037)
- Microsoft SQL Server 2008 Express Edition Service Pack 1 (http://go.microsoft.com/fwlink/?LinkId=166503)
- Windows Identity Foundation for Windows Server 2008 (http://go.microsoft.com/fwlink/?LinkID=160381)

- Windows Identity Foundation for Windows Server 2008 R2 (http://go.microsoft.com/fwlink/?LinkID=166363)
- Microsoft Sync Framework v1.0 (http://go.microsoft.com/fwlink/?LinkID=141237)
- Microsoft Office 2010 Filter Packs (http://go.microsoft.com/fwlink/?LinkId=191851)
- Microsoft Chart Controls for Microsoft .NET Framework 3.5 (http://go.microsoft.com/fwlink/?LinkID=141512)
- Windows PowerShell 2.0 (http://go.microsoft.com/fwlink/?LinkId=161023)
- Microsoft SQL Server 2008 Native Client (http://go.microsoft.com/fwlink/?LinkId=166505)
- Microsoft SQL Server 2008 Analysis Services ADOMD.NET (http://go.microsoft.com/fwlink/?linkid=160390)
- KB979917 - QFE for Sharepoint issues - Perf Counter fix & User Impersonation (http://go.microsoft.com/fwlink/?LinkId=192577)
    - For Windows Server 2008 SP2, download the Windows6.0-KB979917-x64.msu (Vista) file.
    - For Windows Server 2008 R2, download the Windows6.1-KB979917-x64.msu (Win7) file.
- ADO.NET Data Services Update for .NET Framework 3.5 SP1 (http://go.microsoft.com/fwlink/?LinkId=163519) for Windows Server 2008 SP2
- ADO.NET Data Services Update for .NET Framework 3.5 SP1 (http://go.microsoft.com/fwlink/?LinkId=163524) for Windows Server 2008 R2 or Windows 7
- Microsoft Silverlight 3 (http://go.microsoft.com/fwlink/?LinkId=166506)
- Microsoft Office 2010 (http://go.microsoft.com/fwlink/?LinkID=195843)
- SQL Server 2008 R2 Reporting Services Add-in for Microsoft SharePoint Technologies 2010 (http://go.microsoft.com/fwlink/?LinkId=192588)
- SQL Server Remote BLOB Store installation package from the Feature Pack for Microsoft SQL Server 2008 R2. For the download, go to the Download Center (http://go.microsoft.com/fwlink/?LinkID=177388).
- Microsoft Server Speech Platform (http://go.microsoft.com/fwlink/?LinkID=179612)
- Speech recognition language for English (http://go.microsoft.com/fwlink/?LinkID=179613)
- Speech recognition language for Spanish (http://go.microsoft.com/fwlink/?LinkID=179614)
- Speech recognition language for German (http://go.microsoft.com/fwlink/?LinkID=179615)
- Speech recognition language for French (http://go.microsoft.com/fwlink/?LinkID=179616)
- Speech recognition language for Japanese (http://go.microsoft.com/fwlink/?LinkID=179617)
- Speech recognition language for Chinese (http://go.microsoft.com/fwlink/?LinkID=179618)
- Office Communicator 2007 R2 (http://go.microsoft.com/fwlink/?LinkId=196930)
- Microsoft SharePoint Designer 2010 (32-bit) (http://go.microsoft.com/fwlink/?LinkId=196931)
- Microsoft SharePoint Designer 2010 (64-bit) (http://go.microsoft.com/fwlink/?LinkId=196932)

# Plan browser support (SharePoint Server 2010)

Microsoft SharePoint Server 2010 supports several commonly used Web browsers. This article describes different levels of Web browser support, browser compatibility for published sites, and it explains how ActiveX controls affect features.

In this article:

- About planning browser support
- Key planning phase of browser support
- ActiveX controls

## About planning browser support

SharePoint Server 2010 supports several commonly used Web browsers. However, certain Web browsers might cause some SharePoint Server 2010 functionality to be downgraded, limited, or available only through alternative steps. In some cases, functionality might be unavailable for noncritical administrative tasks.

As part of planning your deployment of SharePoint Server 2010, we recommend that you review the browsers used in your organization to ensure optimal performance with SharePoint Server 2010.

If you are using Microsoft Project Server 2010 in your SharePoint Server 2010 farm, pay particular attention to the differences in browser requirements. For more information, see Plan browser support (Project Server 2010).

## Key planning phase of browser support

Browser support is an important part of your SharePoint Server 2010 implementation. Before you install SharePoint Server 2010, ensure that you know which browsers SharePoint Server 2010 supports. The information in this topic covers the following areas:

- Browser support levels
- Browser support matrix
- Browser details
- Browser compatibility for publishing sites

### Browser support levels

Browser support for SharePoint Server 2010 can be divided into three different levels, as follows:

- Supported

A supported Web browser is a Web browser that is supported to work with SharePoint Server 2010, and all features and functionality work. If you encounter any issues, support can help you to resolve these issues.

- Supported with known limitations

  A supported Web browser with known limitations is a Web browser that is supported to work with SharePoint Server 2010, although there are some known limitations. Most features and functionality work, but if there is a feature or functionality that does not work or is disabled by design, documentation on how to resolve these issues is readily available.

- Not tested

  A Web browser that is not tested means that its compatibility with SharePoint Server 2010 is untested, and there may be issues with using the particular Web browser. SharePoint Server 2010 works best with up-to-date, standards-based Web browsers.

## Browser support matrix

The following table summarizes the support levels of commonly used browsers.

| Browser | Supported | Supported with limitations | Not tested |
|---|---|---|---|
| Internet Explorer 8 (32-bit) | X | | |
| Internet Explorer 7 (32-bit) | X | | |
| Internet Explorer 8 (64-bit) | | X | |
| Internet Explorer 7 (64-bit) | | X | |
| Internet Explorer 6 (32-bit) | | | X |
| Mozilla Firefox 3.6 (on Windows operating systems) | | X | |
| Mozilla Firefox 3.6 (on non-Windows operating systems) | | X | |
| Safari 4.04 (on non-Windows operating | | X | |

| Browser | Supported | Supported with limitations | Not tested |
|---|---|---|---|
| systems) | | | |

## Browser details

You should review the details of the Web browser that you have or plan to use in your organization to ensure that the Web browser works with SharePoint Server 2010 and according to your business needs.

**Internet Explorer 8 (32-bit)**

Internet Explorer 8 (32-bit) is supported on the following operating systems:

- Windows Server 2008 R2
- Windows Server 2008
- Windows Server 2003
- Windows 7
- Windows Vista
- Windows XP

**Known limitations**

There are no known limitations for Internet Explorer 8 (32-bit).

**Internet Explorer 7 (32-bit)**

Internet Explorer 7 (32-bit) is supported on the following operating systems:

- Windows Server 2008
- Windows Server 2003
- Windows Vista
- Windows XP

**Known limitations**

There are no known limitations for Internet Explorer 7 (32-bit).

**Internet Explorer 6 (32-bit)**

SharePoint Server 2010 does not support Internet Explorer 6 (32-bit). If you use publishing sites, see Browser compatibility for publishing sites in this article.

**Internet Explorer 8 (64-bit)**

Internet Explorer 8 (64-bit) is supported on the following operating systems:

- Windows Server 2008 R2
- Windows Server 2008
- Windows Server 2003

- Windows 7
- Windows Vista
- Windows XP

**Known limitations**

The following table lists features and their know limitations in Internet Explorer 8 (64-bit).

| Feature | Limitation |
|---|---|
| Connect to Outlook, Connect to Office, and Sync to SharePoint Workspace | Works with an ActiveX control and the stssync:// protocol. Therefore, functionality may be limited without an ActiveX control, such as the one that is included in Microsoft Office 2010. The feature also requires an application that is compatible with the stssync:// protocol, such as Microsoft Outlook. |
| Datasheet view | Requires a 64-bit ActiveX control. Microsoft Office 2010 does not provide a 64-bit version of this control. |
| Edit in Microsoft Office application | Requires a 64-bit ActiveX control. Microsoft Office 2010 does not provide a 64-bit version of this control. |
| Explorer view | Removed in SharePoint Server 2010. Libraries that have been upgraded from earlier versions of SharePoint Server 2010 may still have Explorer views and these may not work. |
| Export to Excel | Downloads a file with an .iqy extension to the Web browser. If Microsoft Excel is not installed, and if no other application is configured to open this file, then this feature will not work. |
| File upload and copy | Requires a 64-bit ActiveX control. Microsoft Office 2010 does not provide a 64-bit version of this control. |
| Microsoft InfoPath 2010 integration | Requires a 64-bit ActiveX control. Microsoft Office 2010 does not provide a 64-bit version of this control. |
| Microsoft PowerPoint 2010 Picture Library integration | Requires a 64-bit ActiveX control, such as the one that is delivered in Microsoft Office 2010. The user can use the following workarounds when no |

| Feature | Limitation |
|---|---|
| | control has been installed:<br>• If a user wants to upload multiple pictures in a picture library, the user must upload one picture at a time by using Upload.aspx.<br>• If a user wants to edit a picture in a picture library, the user must download the picture, edit it, and then upload the picture to the picture library.<br>• If a user wants to download more than one picture from a picture library, the user must download one picture at a time by clicking on the picture link. |
| Microsoft Visio 2010 diagram creation | Requires a 64-bit ActiveX control. Microsoft Office 2010 does not provide a 64-bit version of this control. |
| New Document | Requires a 64-bit ActiveX control. Microsoft Office 2010 does not provide a 64-bit version of this control. Although the **New Document** command may not work, you can use the Upload Document functionality. If you install and configure Office Web Applications on the server, the **New Document** command works, and you can create an Office document in your browser. |
| Send To | Can leverage a 64-bit ActiveX control. Microsoft Office 2010 does not provide a 64-bit version of this control. Without the control, files cannot be sent from one SharePoint farm to another SharePoint farm. However, files can still be sent from one site to another site. |
| Signing Forms (InfoPath Form Services) | Requires a 64-bit ActiveX control. Microsoft Office 2010 does not provide a 64-bit version of this control. |
| Spreadsheet and Database integration | Require a 64-bit ActiveX control. Microsoft Office 2010 does not provide a 64-bit version of this control. The user can use the following workarounds when no control has been installed:<br>• If a user wants to edit a document, the user |

| Feature | Limitation |
|---|---|
| | must download the document, edit it, and then save it back to the server.<br>- In a list that requires a document to be checked out for editing, a user must use the **Edit** menu to check out the document, edit it, and then check it in by using the **Edit** menu.<br>- Export to spreadsheet. Users can export a SharePoint list as a spreadsheet by clicking **Export to Spreadsheet** on the **List** tab on the ribbon. |
| Web Part to Web Part Connections | May require deactivation of browsers pop-up blockers for SharePoint sites. |
| Slide library and PowerPoint 2010 integration | Require a 64-bit ActiveX control. The user can use the following workarounds when no control has been installed:<br>- Delete a slide. Users can delete a slide by first clicking the slide, and then clicking **Delete Slide**. Repeat for each slide. |

**Internet Explorer 7 (64-bit)**

Internet Explorer 7 (64-bit) is supported on the following operating systems:

- Windows Server 2008
- Windows Server 2003
- Windows Vista
- Windows XP

**Known limitations**

The following table lists features and their know limitations in Internet Explorer 7 (64-bit).

| Feature | Limitation |
|---|---|
| Connect to Outlook, Connect to Office, and Sync to SharePoint Workspace | Works with an ActiveX control and the stssync:// protocol. Therefore, functionality may be limited without an ActiveX control, such as the one that is included in Microsoft Office 2010. This feature requires an application that is compatible with the stssync:// protocol, such as Microsoft Outlook. |

| Feature | Limitation |
|---|---|
| Datasheet view | Requires a 64-bit ActiveX control. Microsoft Office 2010 does not provide a 64-bit version of this control. |
| Edit in Microsoft Office application | Requires a 64-bit ActiveX control. Microsoft Office 2010 does not provide a 64-bit version of this control. |
| Explorer view | Removed in SharePoint Server 2010. Libraries that have been upgraded from earlier versions of SharePoint Server 2010 may still have Explorer views. |
| Export to Excel | Downloads a file with an .iqy extension to the Web browser. If Microsoft Excel is not installed, and if no other application is configured to open this file, then this feature will not work. |
| File upload and copy | Requires a 64-bit ActiveX control. Microsoft Office 2010 does not provide a 64-bit version of this control. |
| Microsoft InfoPath 2010 integration | Requires a 64-bit ActiveX control. Microsoft Office 2010 does not provide a 64-bit version of this control. |
| Microsoft PowerPoint 2010 Picture Library integration | Requires a 64-bit ActiveX control, such as the one that is delivered in Microsoft Office 2010. The user can use the following workarounds when no control has been installed:<br>• If a user wants to upload multiple pictures in a picture library, the user must upload one picture at a time by using Upload.aspx.<br>• If a user wants to edit a picture in a picture library, the user must download the picture, edit it, and then upload the picture to the picture library.<br>• If a user wants to download more than one picture from a picture library, the user must download one picture at a time by clicking on the picture link. |
| Microsoft Visio 2010 diagram creation | Requires a 64-bit ActiveX control. Microsoft Office |

| Feature | Limitation |
|---|---|
| | 2010 does not provide a 64-bit version of this control. |
| New Document | Requires a 64-bit ActiveX control. Microsoft Office 2010 does not provide a 64-bit version of this control. Although the **New Document** command may not work, you can use the Upload Document functionality. If you install and configure Office Web Applications on the server, the **New Document** command works, and you can create an Office document in your browser. |
| Send To | Can leverage a 64-bit ActiveX control. Microsoft Office 2010 does not provide a 64-bit version of this control. Without the control, files cannot be sent from one SharePoint farm to another SharePoint farm. However, files can still be sent from one site to another site. |
| Signing Forms (InfoPath Form Services) | Requires a 64-bit ActiveX control. Microsoft Office 2010 does not provide a 64-bit version of this control. |
| Spreadsheet and Database integration | Require a 64-bit ActiveX control. Microsoft Office 2010 does not provide a 64-bit version of this control. The user can use the following workarounds when no control has been installed:<br>• If a user wants to edit a document, the user must download the document, edit it, and then save it back to the server.<br>• In a list that requires a document to be checked out for editing, a user must use the **Edit** menu to check out the document, edit it, and then check it in by using the **Edit** menu.<br>• Export to spreadsheet. Users can export a SharePoint list as a spreadsheet by clicking **Export to Spreadsheet** on the **List** tab on the ribbon. |
| Web Part to Web Part Connections | May require deactivation of browsers pop-up blockers for SharePoint sites. |
| Slide library and PowerPoint 2010 integration | Require a 64-bit ActiveX control. The user can use |

| Feature | Limitation |
|---|---|
| | the following workarounds when no control has been installed:<br>• Delete a slide. Users can delete a slide by first clicking the slide, and then clicking **Delete Slide**. Repeat for each slide. |

### Mozilla Firefox 3.6 (on Windows operating systems)

Mozilla Firefox 3.6 is supported on the following operating systems:

- Windows Server 2008 R2
- Windows Server 2008
- Windows Server 2003
- Windows 7
- Windows Vista
- Windows XP

#### Known limitations

The following table lists features and their know limitations in Mozilla Firefox 3.6 (on Windows operating systems).

| Feature | Limitation |
|---|---|
| Connect to Outlook, Connect to Office, and Sync to SharePoint Workspace | Works with an ActiveX control, but requires a Firefox control adaptor. Microsoft Office 2010 does not provide a Firefox control adaptor for this control. The feature also requires an application that is compatible with the stssync:// protocol, such as Microsoft Outlook. |
| Datasheet view | Requires an ActiveX control, such as the one that is delivered in Microsoft Office 2010, and a Firefox control adaptor. Microsoft Office 2010 does not provide a Firefox control adaptor for this control. |
| Drag and Drop Web Parts | Cannot be moved by using drag and drop on Web Part pages. Users must click **Edit** on the Web Part, select **Modify Web Part**, and then select the zone from the **Layout** section of the Web Part properties page. Web Parts can be moved using drag and drop on Pages. |

| Feature | Limitation |
|---|---|
| Edit in Microsoft Office application | Requires an ActiveX control, such as the one that is delivered in SharePoint Server 2010, and a Firefox control adaptor. For more information about Microsoft Office 2010 Firefox Plug-in, see FFWinPlugin Plug-in (http://go.microsoft.com/fwlink/?LinkId=199867). If you install and configure the Office Web Applications on the server, the Edit functionality works and you can modify Office documents in your browser. This functionality only works with Microsoft Office 2010 or an equivalent product together with a Firefox plug-in. |
| Explorer view | Removed in SharePoint Server 2010. Libraries that have been upgraded from earlier versions of SharePoint Server 2010 may still have Explorer views, and these may not work. Explorer view requires Internet Explorer. |
| Export to Excel | Downloads a file with an .iqy extension to the Web browser. If Microsoft Excel is not installed, and if no other application is configured to open this file, then this feature will not work. |
| File upload and copy | Requires an ActiveX control, such as the one that is delivered in Microsoft Office 2010, and a Firefox control adaptor. Microsoft Office 2010 does not provide a Firefox control adaptor for this control. |
| Microsoft InfoPath 2010 integration | Requires an ActiveX control, such as the one that is delivered in Microsoft Office 2010, and a Firefox control adaptor. Microsoft Office 2010 does not provide a Firefox control adaptor for this control. |
| Microsoft PowerPoint 2010 Picture Library integration | Requires an ActiveX control, such as the one that is delivered in Microsoft Office 2010, and a Firefox control adaptor. Microsoft Office 2010 does not provide a Firefox control adaptor for this control. The user can use the following workarounds when no control has been installed:<br>• If a user wants to upload multiple pictures in a picture library, the user must upload one |

| Feature | Limitation |
|---|---|
| | picture at a time by using Upload.aspx.<br>• If a user wants to edit a picture in a picture library, the user must download the picture, edit it, and then upload the picture to the picture library.<br>• If a user wants to download more than one picture from a picture library, the user must download one picture at a time by clicking on the picture link. |
| Microsoft Visio 2010 diagram creation | Requires an ActiveX control, such as the one delivered in Microsoft Office 2010, and a Firefox control adaptor. Microsoft Office 2010 does not provide a Firefox control adaptor for this control. |
| New Document | Requires an ActiveX control, such as the one delivered in Microsoft Office 2010, and a Firefox control adaptor. For more information about Microsoft Office 2010 Firefox Plug-in, see FFWinPlugin Plug-in (http://go.microsoft.com/fwlink/?LinkId=199867). Although the **New Document** command may not work, you can use the Upload Document functionality. If you install and configure Office Web Applications on the server, the **New Document** command works, and you can create an Office document in your browser. |
| Rich Text Editor – Basic Toolbar | A user can update the Rich Text Editor basic toolbar to a Full Rich Text Editor that includes the ribbon by changing the field's properties, as follows: On the FldEdit.aspx, in the **List Settings** menu, select **Specific Field Settings**. Next, under **Columns**, click **Description**. In the **Additional Columns Settings** section, under **Specify the type of text to allow**, select **Enhanced rich text (Rich text with pictures, tables, and hyperlinks)**. |
| Send To | Can leverage an ActiveX control, such as the one that is delivered in Microsoft Office 2010, and a Firefox control adaptor. Microsoft Office 2010 does |

| Feature | Limitation |
|---|---|
| | not provide a Firefox control adaptor for this control. Without the control, files cannot be sent from one SharePoint farm to another SharePoint farm. However, files can still be sent from one site to another site. |
| Signing Forms (InfoPath Form Services) | Requires an ActiveX control, such as the one that is delivered in Microsoft Office 2010, and a Firefox control adaptor. Microsoft Office 2010 does not provide a Firefox control adaptor for this control. |
| Spreadsheet and Database integration | Require ActiveX controls, such as those that are delivered in Microsoft Office 2010, and Firefox control adaptors. Microsoft Office 2010 does not provide a Firefox control adaptor for this control. The user can use the following workarounds when no control has been installed:<br>• If a user wants to edit a document, the user must download the document, edit it, and then save it back to the server.<br>• In a list that requires a document to be checked out for editing, a user must use the **Edit** menu to check out the document, edit it, and then check it in by using the **Edit** menu.<br>• Export to spreadsheet. Users can export a SharePoint list as a spreadsheet by clicking **Export to Spreadsheet** on the **List** tab on the ribbon. |
| Web Part to Web Part Connections | May require deactivation of browsers pop-up blockers for SharePoint sites. |
| Slide library and PowerPoint 2010 integration | Require ActiveX controls, such as those that are delivered in Microsoft Office 2010, and Firefox control adaptors. Microsoft Office 2010 does not provide a Firefox control adaptor for this control. The user can use the following workarounds when no control has been installed:<br>• Delete a slide. Users can delete a slide by first clicking the slide, and then clicking **Delete Slide**. Repeat for each slide. |

| Feature | Limitation |
|---|---|
| | The following features do not work on this platform:<br>• Copy a slide to a presentation. This feature enables users to add a slide to a PowerPoint 2010 presentation.<br>• Publish a slide. This feature enables users to upload a single slide from a PowerPoint 2010 presentation to a slide library. Microsoft Office must be installed on the client computer. |

**Mozilla FireFox 3.6 (on non-Windows operating systems)**

Mozilla FireFox 3.6 is supported on the following operating systems:

- Mac OSX
- UNIX/Linux

**Known limitations**

The following table lists features and their know limitations in Mozilla FireFox 3.6 (on non-Windows operating systems).

| Feature | Limitation |
|---|---|
| Connect to Outlook, Connect to Office, and Sync to SharePoint Workspace | Requires an application that is compatible with the stssync:// protocol, such as Microsoft Outlook. |
| Datasheet view | Requires an ActiveX control that is not supported on this platform. Microsoft Office 2010 does not provide a Firefox control adaptor for this control. |
| Drag and Drop Web Parts | Cannot be moved by using drag and drop on Web Part pages. Users must click **Edit** on the Web Part, select **Modify Web Part**, and then select the zone from the **Layout** section of the Web Part properties page. Web Parts can be moved using drag and drop on Pages. |
| Edit in Microsoft Office application | Requires an ActiveX control that is not supported on this platform. If you install and configure the Office Web Applications on the server, the Edit functionality works and you can modify Office documents in your browser. |

| Feature | Limitation |
|---|---|
| Explorer view | Removed in SharePoint Server 2010. Libraries that have been upgraded from earlier versions of SharePoint Server 2010 may still have Explorer views, and these may not work. Explorer view requires Internet Explorer. |
| Export to Excel | Downloads a file with an .iqy extension to the Web browser. Requires an application that is configured to open this file. |
| File upload and copy | Requires an ActiveX control that is not support on this platform. |
| Microsoft InfoPath 2010 integration | Requires an ActiveX control that is not support on this platform. |
| Microsoft PowerPoint 2010 Picture Library integration | Requires an ActiveX control that is not supported on this platform. Microsoft Office 2010 does not provide a Firefox control adaptor for this control. The user can use the following workarounds when no control has been installed:<br>• If a user wants to upload multiple pictures in a picture library, the user must upload one picture at a time by using Upload.aspx.<br>• If a user wants to edit a picture in a picture library, the user must download the picture, edit it, and then upload the picture to the picture library.<br>• If a user wants to download more than one picture from a picture library, the user must download one picture at a time by clicking on the picture link. |
| Microsoft Visio 2010 diagram creation | Requires an ActiveX control that is not supported on this platform. |
| New Document | Requires an ActiveX control that is not supported on this platform. Although the **New Document** command may not work, you can use the Upload Document functionality. If you install and configure Office Web Applications on the server, the **New Document** command works, and you can create |

| Feature | Limitation |
|---|---|
| | an Office document in your browser. |
| Rich Text Editor – Basic Toolbar | A user can update the Rich Text Editor basic toolbar to a Full Rich Text Editor that includes the ribbon by changing the field's properties, as follows: On the FldEdit.aspx, in the **List Settings** menu, select **Specific Field Settings**. Next, under **Columns**, click **Description**. In the **Additional Columns Settings** section, under **Specify the type of text to allow**, select **Enhanced rich text (Rich text with pictures, tables, and hyperlinks)**. |
| Send To | Can leverage an ActiveX control that is not supported on this platform. Without the control, files cannot be sent from one SharePoint farm to another SharePoint farm. However, files can still be sent from one site to another site. |
| Signing Forms (InfoPath Form Services) | Requires an ActiveX control that is not supported on this platform. |
| Spreadsheet and Database integration | Require ActiveX controls that is not supported on this platform. The user can use the following workarounds when no control has been installed:<br>• If a user wants to edit a document, the user must download the document, edit it, and then save it back to the server.<br>• In a list that requires a document to be checked out for editing, a user must use the **Edit** menu to check out the document, edit it, and then check it in by using the **Edit** menu.<br>• Export to spreadsheet. Users can export a SharePoint list as a spreadsheet by clicking **Export to Spreadsheet** on the **List** tab on the ribbon. |
| Web Part to Web Part Connections | May require deactivation of browsers pop-up blockers for SharePoint sites. |
| Slide library and PowerPoint 2010 integration | Require ActiveX controls that is not supported on this platform. The user can use the following workarounds when no control has been installed: |

| Feature | Limitation |
|---|---|
| | • Delete a slide. Users can delete a slide by first clicking the slide, and then clicking **Delete Slide**. Repeat for each slide.<br>The following features do not work on this platform:<br>• Copy a slide to a presentation. This feature enables users to add a slide to a PowerPoint 2010 presentation.<br>• Publish a slide. This feature enables users to upload a single slide from a PowerPoint 2010 presentation to a slide library. Microsoft Office must be installed on the client computer. |

**Note:**
FireFox browsers on UNIX/Linux systems may not work with the Web Part menu.

**Note:**
Some ActiveX features, such as list Datasheet view and the control that displays user presence information, do not work in Mozilla Firefox 3.6. Firefox users can use the Microsoft Office 2010 Firefox Plug-in to launch documents.

**Safari 4.04 (on non-Windows operating systems)**

Safari 4.0.4 is supported on the following operating systems:

- Mac OSX (Version 10.6, Snow Leopard)

**Known limitations**

The following table lists features and their know limitations in Safari 4.04 (on non-Windows operating systems).

| Feature | Limitation |
|---|---|
| Connect to Outlook, Connect to Office, and Sync to SharePoint Workspace | Requires an application that is compatible with the stssync:// protocol, such as Microsoft Outlook. |
| Datasheet view | Requires an ActiveX control that is not supported on this platform. |
| Drag and Drop Web Parts | Cannot be moved by using drag and drop on Web Part pages. Users must click **Edit** on the Web Part, select **Modify Web Part**, and then select the zone from the **Layout** section of the Web Part |

| Feature | Limitation |
|---|---|
| | properties page. Web Parts can be moved using drag and drop on Pages. |
| Edit in Microsoft Office application | Requires an ActiveX control that is not supported on this platform. If you install and configure the Office Web Applications on the server, the Edit functionality works and you can modify Office documents in your browser. |
| Explorer view | Removed in SharePoint Server 2010. Libraries that have been upgraded from earlier versions of SharePoint Server 2010 may still have Explorer views. Explorer view requires Internet Explorer. |
| Export to Excel | Downloads a file with an .iqy extension to the Web browser. Requires an application that is configured to open this file. |
| File upload and copy | Requires an ActiveX control that is not supported on this platform. |
| Microsoft InfoPath 2010 integration | Requires an ActiveX control that is not supported on this platform. |
| Microsoft PowerPoint 2010 Picture Library integration | Requires an ActiveX control that is not supported on this platform. The user can use the following workarounds when no control has been installed:<br>• If a user wants to upload multiple pictures in a picture library, the user must upload one picture at a time by using Upload.aspx.<br>• If a user wants to edit a picture in a picture library, the user must download the picture, edit it, and then upload the picture to the picture library.<br>• If a user wants to download more than one picture from a picture library, the user must download one picture at a time by clicking on the picture link. |
| Microsoft Visio 2010 diagram creation | Requires an ActiveX control that is not supported on this platform. |
| New Document | Requires an ActiveX control that is not supported on this platform. Although the **New Document** |

| Feature | Limitation |
|---|---|
| | command may not work, you can use the Upload Document functionality. If you install and configure Office Web Applications on the server, the **New Document** command works, and you can create an Office document in your browser. |
| Rich Text Editor – Basic Toolbar | A user can update the Rich Text Editor basic toolbar to a Full Rich Text Editor that includes the ribbon by changing the field's properties, as follows: On the FldEdit.aspx, in the **List Settings** menu, select **Specific Field Settings**. Next, under **Columns**, click **Description**. In the **Additional Columns Settings** section, under **Specify the type of text to allow**, select **Enhanced rich text (Rich text with pictures, tables, and hyperlinks)**. |
| Send To | Can leverage an ActiveX control that is not supported on this platform. Without the control, files cannot be sent from one SharePoint farm to another SharePoint farm. However, files can still be sent from one site to another site. |
| Signing Forms (InfoPath Form Services) | Requires an ActiveX control that is not supported on this platform. |
| Spreadsheet and Database integration | Require ActiveX controls that are not supported on this platform. The user can use the following workarounds when no control has been installed:<br>• If a user wants to edit a document, the user must download the document, edit it, and then save it back to the server.<br>• In a list that requires a document to be checked out for editing, a user must use the **Edit** menu to check out the document, edit it, and then check it in by using the **Edit** menu.<br>• Export to spreadsheet. Users can export a SharePoint list as a spreadsheet by clicking **Export to Spreadsheet** on the **List** tab on the ribbon. |
| Web Part to Web Part Connections | May require deactivation of browsers pop-up |

| Feature | Limitation |
|---|---|
|  | blockers for SharePoint sites. |
| Slide library and PowerPoint 2010 integration | Require ActiveX controls that is not supported on this platform. The user can use the following workarounds when no control has been installed:<br>• Delete a slide. Users can delete a slide by first clicking the slide, and then clicking **Delete Slide**. Repeat for each slide.<br>The following features do not work on this platform:<br>• Copy a slide to a presentation. This feature enables users to add a slide to a PowerPoint 2010 presentation.<br>• Publish a slide. This feature enables users to upload a single slide from a PowerPoint 2010 presentation to a slide library. Microsoft Office must be installed on the client computer. |

## Browser compatibility for publishing sites

For publishing sites, the Web Content Management features built into SharePoint Server 2010 provide a deep level of control over the markup and styling of the reader experience. Page designers can use these features to help ensure that the pages they design are compatible with additional browsers, including Internet Explorer 6, for viewing content. However, it is the page designers' responsibility to create pages that are compatible with the browsers that they want to support.

A standards-based browser, such as Internet Explorer 8 or Firefox 3.x, is required to author content.

# ActiveX controls

Some of the features in SharePoint Server 2010 use ActiveX controls. In secure environments, these controls must be able to work on the client computer before their features will function. Some ActiveX controls, such as those included in Microsoft Office 2010, does not work with 64-bit browser versions. For Microsoft Office 2010 (64-bit), only the following controls work with 64-bit browsers:

- ppslax.dll – Slide library and PowerPoint 2010 integration
- name.dll – Presence information

# URL path length restrictions (SharePoint Server 2010)

This article discusses the specific URL path length and character restrictions in Microsoft SharePoint Server 2010, Internet Explorer 7, and Internet Explorer 8 that you should be aware of when planning sites, navigation, and structure. This article does not discuss URL length limitations in other browsers. For this information, see the browser documentation.

In this article:

- Understanding URL and path lengths
- URL path length limitations
- Resolving URL length problems

# Understanding URL and path lengths

This section discusses URL composition, how SharePoint Server 2010 builds URLs, how URLs are encoded and lengthened, and passed as parameters in other URLs.

## SharePoint URL composition

The total length of a SharePoint URL equals the length of the folder or file path, including the protocol and server name and the folder or file name, plus any parameters that are included as part of the URL. The formula is as follows:

URL = protocol + server name + folder or file path + folder or file name+ parameters

For example, the following is a typical URL path to a file stored in Microsoft SharePoint Server 2010:

*http://www.contoso.com/sites/marketing/documents/Shared%20Documents/Promotion/Some%20File.xlsx*

Where the parts of the URL path are as listed in the following table.

| URL part | Example |
|---|---|
| Protocol | http:// |
| Server name | www.contoso.com/ |
| Folder or file path | sites/marketing/documents/Shared%20Documents/Promotion/ |
| File name | Some%20File.xlsx |

When you go to the site and open the file with Microsoft Office Web Apps, the URL will be as follows:

*http://www.contoso.com/sites/marketing/documents/_layouts/xlviewer.aspx?id=/sites/marketing/documents/Shared%20Documents/Promotion/Some%20File.xlsx&Source=http%3A%2F%2Fwww%2Econtoso%2Ecom%2Fsites%2Fmarketing%2Fdocuments%2FShared%2520Documents%2FForms%2FAllItems%2Easpx%3FRootFolder%3D%252Fsites%252Fmarketing%252Fdocuments%252FShared%2520Documents%252FPromotion%26FolderCTID%3D0x012000F2A09653197F4F4F919923797C42ADEC&DefaultItemOpen=1*

Where the parts of the URL path are as listed in the following table.

| URL part | Example |
|---|---|
| Protocol | http:// |
| Server name | www.contoso.com/ |
| Folder or file path | sites/marketing/documents/Shared%20Documents/Promotion/ |
| Folder or file name | xlviewer.aspx |
| Parameters | ?id=/sites/marketing/documents/Shared%20Documents/Promotion/Some%20File.xlsx<br>&Source=http%3A%2F%2Fwww%2Econtoso%2Ecom%2Fsites%2Fmarketing%2Fdocuments%2FShared%2520Documents%2FForms%2FAllItems%2Easpx%3FRootFolder%3D%252Fsites%252Fmarketing%252Fdocuments%252FShared%2520Documents%252FPromotion%26FolderCTID%3D0x012000F2A09653197F4F4F919923797C42ADEC<br>&DefaultItemOpen=1 |

## URL Encoding

URL encoding ensures that all browsers will correctly transmit text in URL strings. Characters such as a question marks (?), ampersands (&), slash marks (/), and spaces might be truncated or corrupted by some browsers. SharePoint Server 2010 adheres to the standards for URL encoding that are defined in

The Internet Engineering Task Force (IETF) RFC 3986 (http://go.microsoft.com/fwlink/?LinkId=195564&clcid=0x409).

In the URL path example earlier in this article, the Source parameter contains a double-encoded path and is 262 characters. The first de-encoding reveals:

*&Source=http://www.contoso.com/sites/marketing/documents/Shared%20Documents/Forms/AllItems.aspx?RootFolder=%2Fsites%2Fmarketing%2Fdocuments%2FShared%20Documents%2FPromotion&FolderCTID=0x012000F2A09653197F4F4F919923797C42ADEC* which is 216 characters.

De-encoded again reveals:

*&Source=http://www.contoso.com/sites/marketing/documents/Shared Documents/Forms/AllItems.aspx?RootFolder=/sites/marketing/documents/Shared Documents/Promotion&FolderCTID=0x012000F2A09653197F4F4F919923797C42ADEC* which is 200 characters.

If you have non-standard ASCII characters, such as high-ASCII or double-byte Unicode characters, in the SharePoint URL, each of those characters is URL-encoded into two or more ASCII characters when they are passed to the Web browser. Thus, a URL with many high-ASCII characters or double-byte Unicode characters can become longer than the original un-encoded URL. The list below gives examples of the multiplication factors:

- High-ASCII characters — for example, (!, ", #, $, %, &, [Space]): multiplication factor = 3
- Double byte Unicode characters — for example, Japanese, Chinese, Korean, Hindi: multiplication factor = 9

For example, when you translate the names of sites, library, folder, and file in the URL path *http://www.contoso.com/sites/marketing/documents/Shared%20Documents/Promotion/Some%20File.xlsx* into Japanese, the resulted encoded URL path will become something like the following:

*http://www.contoso.com/sites/%E3%83%9E%E3%83%BC%E3%82%B1%E3%83%86%E3%82%A3%E3%83%B3%E3%82%B0/%E6%96%87%E6%9B%B8/DocLib/%E3%83%97%E3%83%AD%E3%83%A2%E3%83%BC%E3%82%B7%E3%83%A7%E3%83%B3/%E3%83%95%E3%82%A1%E3%82%A4%E3%83%AB.xlsx*. This path is 224 characters, whereas the original URL path is only 94 characters.

**Important:**
The following characters cannot be used in an un-encoded URL: (~, #, %, &, *, {}, \, :, <>, /, +, |, “).

## URL parameters

URL parameters are data that are included as part of the URL that are processed. These parameters are also URL-encoded and can be encoded multiple times, producing very long URLs.

For example, if you browse to a list, the URL might be something like the following: *http://www.contoso.com/sites/marketing/documents/Shared%20Documents/Forms/AllItemA.aspx?RootFolder=%2Fsites%2Fmarketing%2Fdocuments%2FShared%20Documents%2FPFPromoti&FolderCTID=0x012000F2A09653197F4F4F919923797C42ADEC&View={CD527605-9A7A-448D-9A35-67A33EF9F766}*. This URL is 260 characters.

If you then click **Create View** on the **Library** tab, the entire URL is included in the resulting URL as the source parameter and it is encoded to be much longer — for example, *http://www.contoso.com/sites/marketing/documents/_layouts/ViewType.aspx?List=%7BED6E21E0%2DDF28%2D4165%2DBC3E%2D5371987CC2D2%7D&Source=http%3A%2F%2Fwww%2Econtoso%2Ecom%2Fsites%2Fmarketing%2Fdocuments%2FShared%2520Documents%2FForms%2FAllItems%2Easpx%3FRootFolder%3D%252Fsites%252Fmarketing%252Fdocuments%252FShared%2520Documents%252FPromotion%26FolderCTID%3D0x012000F2A09653197F4F4F919923797C42ADEC%26View%3D%7BCD527605%2D9A7A%2D448D%2D9A35%2D67A33EF9F766%7D*. This URL is 457 characters.

**Important**

- SharePoint Server 2010 truncates the URL source parameter if the total URL length to be passed to Internet Explorer is more than 1950 bytes. The source parameter is a reference to a previously visited page. The result of the truncation of the source parameter is that the user will be referred back to default location rather than the location specified in the source parameter.
- Other parameters, such as sort orders, root folder parameters, and views are not truncated.

# URL path length limitations

This section discusses the different URL length limitations in SharePoint Server 2010 and Internet Explorer, and how to plan for URL path lengths.

## SharePoint URL path length limitations

The limitations In this section apply to the total length of the URL path to a folder or a file in SharePoint Server 2010 but not to the length of any parameters. Also, these limitations apply only to un-encoded URLs, not to encoded URLs. There is no limit to encoded URLs in SharePoint Server 2010. The limitations are the following:

- **260 Unicode (UTF-16) code units** – the characters in a full file path, not including a domain/server name.
- **256 Unicode (UTF-16) code units** – the characters in a full folder path, not including the file name and the domain/server name.
- **128 Unicode (UTF-16) code units** - characters in a path component, that is, a file or folder name.
- **260 Unicode (UTF-16) code units** – the characters in a full path, including a domain/server name for use with Office clients.
- **256 Unicode (UTF-16) code units** – the characters in a full path including the domain/server name, for use with Active X controls.

For more information, see Microsoft Knowledge Base article 894630, You receive a "The specified file or folder name is too long" error message (http://go.microsoft.com/fwlink/?LinkId=195567&clcid=0x409).

**Note:**

**Understanding code units** - In most cases, one UTF-16 character equals one UTF-16 code unit. However, characters that use Unicode code points greater than U+10000 will equal two UTF-16 code units. These characters include, but are not limited to, Japanese or Chinese surrogate pair characters. If your paths include these characters, the URL length will exceed the URL length limitation with fewer than 256 or 260 characters.

## Internet Explorer URL length limitations

Internet Explorer also has limitations that are separate from those in SharePoint Server 2010. Even though you make the SharePoint Server 2010 URL path shorter than the limitations, you might experience an Internet Explorer URL length limitation because of added parameters and encoding of the URL. You must use the most restrictive limitation as a guideline for planning URL lengths.

Both Internet Explorer 7 and Internet Explorer 8 have a maximum URL length of 2,083 UTF-8 characters and a maximum path length of 2,048 UTF-8 characters. However, in Internet Explorer 7, under certain circumstances, the effective URL length limitation is 1024 UTF-8 characters, not 2083 UTF-8 characters. For more information about the URL length limits in Internet Explorer, see Microsoft Knowledge Base article 208427, Maximum URL length is 2,083 characters in Internet Explorer (http://go.microsoft.com/fwlink/?LinkId=195568&clcid=0x409).

**Important:**
Unless all of the browsers in the environment are Internet Explorer 8, use the effective limit of 1024 UTF-8 characters.

# Resolving URL length problems

There are several ways that you can resolve or mitigate URL length problems in the SharePoint Server 2010 environment. The following list provides suggestions:

- Upgrade all the end-user browsers to Internet Explorer 8, which has a longer URL length limit.
- Use shorter names for sites, folders, and documents and control the depth of the site and folder structures to reduce the lengths of URLs.
- If possible or allowed, use ASCII names for sites, folders, and documents. This will avoid situations where the URL will be lengthened by being encoded.
- To reduce the risk that the SharePoint Server 2010 end-users will encounter problems because of URL length limitations, we recommend that you apply the following effective limits in the deployment:
  - **256 Unicode (UTF-16) Code units** - the effective file path length limitation, including a domain/server name
  - **128 Unicode (UTF-16) Code units** - the path component length limitation

# IP support (SharePoint Server 2010)

This article explains the support for Internet Protocol Version 4 (IPv4) and Internet Protocol Version 6 (IPv6) addressing in Microsoft SharePoint 2010 Products.

SharePoint 2010 Products support the following environments:

- Pure IPv4 environment
- Mixed IPv4 and IPv6 environment
- Pure IPv6 environment

In a SharePoint environment, "mixed" can be defined as one of the following likely scenarios:

- Both IPv4 and IPv6 protocols are running in your environment.
- Some of your client computers are using IPv4 and some of them are using IPv6.
- Your client computers are using IPv4, but the computer running Microsoft SQL Server is using IPv6.

By default, the IPv6 protocol and the IPv4 protocol are both installed and enabled in Windows Server 2008 and Windows Server 2008 R2. When both IPv4 and IPv6 are enabled, IPv6 is given preference over IPv4. Additionally, you can remove the IPv4 protocol so that the computer runs IPv6 exclusively.

To determine what version is being used, you can use the IPConfig.exe tool. For additional information, see IPConfig (http://go.microsoft.com/fwlink/?LinkId=122336&clcid=0x409).

The following list shows other important considerations regarding IPv6:

- For any computer that is authenticated by using a domain controller and is only running IPv6 within a SharePoint 2010 Products environment, the domain controller must be running Windows Server 2008 or Windows Server 2008 R2. Ensure that you use the correct service pack and any additional software prerequisites. For more information, see Hardware and software requirements (SharePoint Server 2010).
- All versions of Microsoft SQL Server supported for SharePoint 2010 Products also support IPv6. For more information about IPv6 support for SQL Server 2008, see Connecting Using IPv6 (http://go.microsoft.com/fwlink/?LinkId=183115). For more information about IPv6 support for SQL Server 2005, see Connecting Using IPv6 (http://go.microsoft.com/fwlink/?LinkId=183118).
- In SharePoint 2010 Products, when using IPv6 protocol, all end-user Uniform Resource Locators (URLs) must be based on DNS names with AAAA records. Browsing to SharePoint URLs that use IPv6 literal addresses is not supported. An example of a literal address URL is http://[2001:db8:85a3:8d3:1319:8a2e:370:7344]. However, SharePoint 2010 Products support entering IPv6 literal addresses for certain farm administration functionality, such as entering the server name when creating or attaching databases. For server names that use a literal address format, you must enclose the literal address within square brackets. For more information about AAAA records, see Adding a Resource Record to a Forward Lookup Zone (http://go.microsoft.com/fwlink/?LinkId=181956).

For additional information about IPv6, see Internet Protocol Version 6 (IPv6) (http://go.microsoft.com/fwlink/?LinkId=120794&clcid=0x409) and IP Addressing (http://go.microsoft.com/fwlink/?LinkId=120795&clcid=0x409).

### See Also

Internet Protocol, Version 6 (IPv6) Specification

# Windows Server 2008 R2 and SharePoint Server 2010: Better Together (white paper)

This white paper describes the benefits of deploying Microsoft SharePoint Server 2010 on the Windows Server 2008 R2 Enterprise operating system and scenarios to which the features of Windows Server 2008 R2 Enterprise can be applied.

Download this white paper as a .docx file (http://go.microsoft.com/fwlink/?LinkId=199051).

Download this white paper as a PDF file (http://go.microsoft.com/fwlink/?LinkId=199052).

Download this white paper as an XPS file (http://go.microsoft.com/fwlink/?LinkId=199053).

### See Also

Business Productivity at Its Best: Microsoft Office 2010 and SharePoint Server 2010 Better Together (white paper)

SQL Server 2008 R2 and SharePoint 2010 Products: Better Together (white paper) (SharePoint Server 2010)

Hardware and software requirements (SharePoint Server 2010)

# SQL Server 2008 R2 and SharePoint 2010 Products: Better Together (white paper) (SharePoint Server 2010)

Choosing an edition of Microsoft SQL Server 2008 R2 is an important step when planning a Microsoft SharePoint Server 2010 deployment. This paper describes the benefits of deploying on SQL Server 2008 R2 Enterprise Edition and scenarios in which its features can be applied.

Download this white paper as a Word document (.docx)
(http://go.microsoft.com/fwlink/?LinkID=187264).

# Business Productivity at Its Best: Microsoft Office 2010 and SharePoint Server 2010 Better Together (white paper)

This white paper shows how Microsoft Office 2010 and Microsoft SharePoint 2010 Products contribute to the powerful architectural design of the Microsoft Business Productivity Infrastructure (BPI). It provides an overview of Office and SharePoint features working together in past versions, and focuses on the integration features of the Office 2010 experience with SharePoint 2010 Products.

Download this white paper as a PDF file: Business Productivity at Its Best (http://go.microsoft.com/?linkid=9690494)

# Logical architecture planning (SharePoint Server 2010)

This section contains articles to help you learn about and plan logical architectures for Microsoft SharePoint Server 2010.

In this section:

- Services architecture planning (SharePoint Server 2010)
- Logical architecture components (SharePoint Server 2010)
- Design sample: Corporate deployment (SharePoint Server 2010)Corporate deployment design sample
- Plan for host-named site collections (SharePoint Server 2010)

# Services architecture planning (SharePoint Server 2010)

This article describes the services architecture for sharing service applications and provides example architectures for Microsoft SharePoint Server 2010.

In this article:

- About service applications
- Services infrastructure and design principles
- Deploying service applications across farms
- Planning considerations for services that access external data sources
- Example architectures
- Single farm, single service group
- Single farm, multiple service groups
- Enterprise services farms
- Specialized service farms
- Cross-organization farms

When planning your services architecture, consider the following questions:

- What service applications does your organization require?
- Do any teams require dedicated service applications?
- How many farms does your organization require?
- Are there opportunities to share services across farms?
- Do the needs of your organization warrant a centralized services farm?

The following poster-size models are also available to use with this article. You can modify the diagrams within the models to represent your own organization plans.

- Services in Microsoft SharePoint 2010 Products
    - Visio (http://go.microsoft.com/fwlink/?LinkID=167090)
    - PDF (http://go.microsoft.com/fwlink/?LinkID=167092)
    - XPS (http://go.microsoft.com/fwlink/?LinkID=167091)
- Cross-farm services in SharePoint 2010 Products
    - Visio(http://go.microsoft.com/fwlink/?LinkID=167093)
    - PDF (http://go.microsoft.com/fwlink/?LinkID=167095)
    - XPS (http://go.microsoft.com/fwlink/?LinkID=167094)

# About service applications

SharePoint Server 2010 includes a set of services that can be shared across Web applications. These services are called *service applications*. Some service applications can be shared across farms. Sharing service applications across Web applications and farms greatly reduces the resources required to provide these services across multiple sites.

The following table lists service applications that are included with SharePoint 2010 Products.

| **Service applications** | **Description** | **SharePoint Foundation 2010** | **SharePoint Server 2010 Standard** | **SharePoint Server 2010 Enterprise** |
|---|---|---|---|---|
| Access Services | Lets users view, edit, and interact with Access 2010 databases in a Web browser. | | | X |
| Business Data Connectivity service | Gives access to line-of-business data systems. | X | X | X |
| Excel Services Application | Lets users view and interact withExcel 2010 files in a Web browser. | | | X |
| Managed Metadata service | Manages taxonomy hierarchies, keywords and social tagging infrastructure, and publish content types across site collections. | | X | X |
| PerformancePoint Service Application | Provides the capabilities of PerformancePoint. | | | |
| Search service | Crawls content, produces index partitions, and serves search queries. | | X | X |
| Secure Store Service | Provides single sign-on authentication to | | X | X |

| | | | | |
|---|---|---|---|---|
| | access multiple applications or services. | | | |
| State service | Provides temporary storage of user session data for SharePoint Server components. | | X | X |
| Usage and Health Data Collection service | Collects farm wide usage and health data, and provides the ability to view various usage and health reports. | X | X | X |
| User Profile service | Adds support for My Site Web sites, profile pages, social tagging and other social computing features. | | X | X |
| Visio Graphics Service | Lets users view and refresh published Visio 2010 diagrams in a Web browser. | | | X |
| Web Analytics service | Provides Web service interfaces. | | X | X |
| Word Automation Services | Performs automated bulk document conversions. | | X | X |
| Microsoft SharePoint Foundation Subscription Settings Service | Provides multi-tenant functionality for service applications. Tracks subscription IDs and settings for services that are deployed in partitioned mode. Deployed through Windows PowerShell only. | X | X | X |

Some services are provided by other Microsoft products, including the services listed in the following table.

| **Service application** | **Description** |
|---|---|
| Office Web Apps services:<br>• Word Viewing Service<br>• PowerPoint Service<br>• Excel Calculation Services | Office Web Apps is a Web-based productivity offering from Microsoft Office 2010 suites. Office Web Apps services include companions to Microsoft Word 2010, Microsoft Excel 2010, Microsoft PowerPoint 2010, and Microsoft OneNote 2010. These Web-based applications are stand-alone applications focused on offering access to Word 2010, PowerPoint 2010, Excel 2010, and OneNote 2010 documents through any browser across multiple platforms, lightweight creation and editing capabilities in standard formats, sharing and collaboration on those documents through the browser, and various Web-enabled scenarios. Documents created by using Office Web Apps are no different from documents that were created by using the corresponding desktop applications. The associated services are used to prepare documents for viewing and editing in a Web browser. |
| Microsoft Project Server 2010 | Microsoft Project Server 2010 hosts one or more Microsoft Project Web Access instances, exposes scheduling functionality and other middle-tier calculations on Microsoft Project data, and exposes Web services for interacting with Microsoft Project 2010 data. |

Service applications are different from the services that are started and stopped on specific servers and listed on the Services on Server page in the SharePoint Central Administration Web site. Some of the services listed on this page are associated with service applications, but service applications represent specific instances of services that can be configured and shared in specific ways.

# Services infrastructure and design principles

SharePoint 2010 Products improves the services infrastructure that was introduced in the previous version. In SharePoint 2010 Products, the infrastructure for hosting services moves into SharePoint Foundation 2010 and the configuration of service offerings is much more flexible. Individual services can be configured independently, and third-party companies can add services to the platform.

Sharing services is no longer exclusive to SharePoint Server, and services are no longer contained in Shared Services Providers (SSPs).

## Deploying services

You deploy service applications within a farm by using one of the following methods:

- Selecting services when you run the SharePoint Products Configuration Wizard.
- Adding services one by one on the Manage Service Applications page in the Central Administration site.
- Using Windows PowerShell.

## More granular configuration of services

The updated services infrastructure gives you more control over which services are deployed and how service applications are shared:

1. You can deploy only the service applications that are needed to a farm.
2. Web applications can be configured to use only the service applications that are needed, instead of all the services that have been deployed.
3. You can deploy multiple instances of the same service in a farm and assign unique names to the resulting service applications.
4. You can share service applications across multiple Web applications within the same farm.

You can choose the service applications for a Web application when you create the Web application. You can also modify the service applications that are associated with a Web application later.

## Service application groups

By default, all service applications are included in a default group, unless you change this setting for a service application when it is created. You can add and remove service applications from the default group at any time.

When you create a Web application, you can select the default group or you can create a custom group of service applications. You create a custom group of service applications by selecting only the service applications that you want the Web application to use.

The following screen shot shows a list of service applications for an example farm that can be selected if **custom** is selected when a Web application is created. Only the first few service applications are included in the picture.

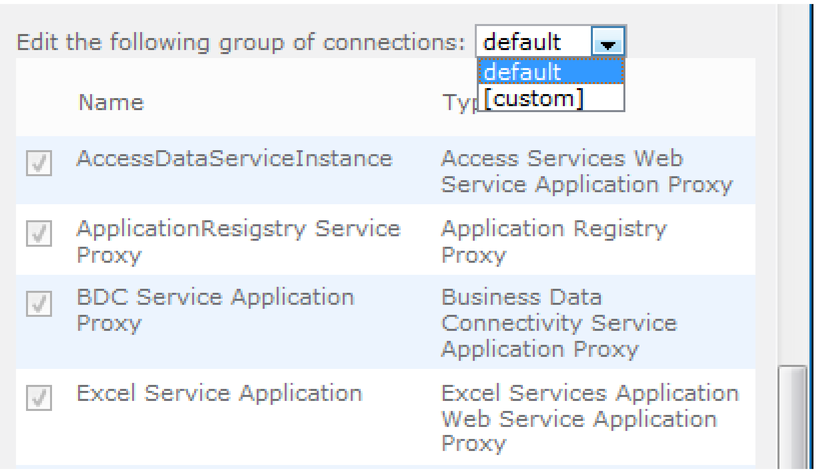

Edit the following group of connections: default

| | Name | Typ |
|---|---|---|
| ☑ | AccessDataServiceInstance | Access Services Web Service Application Proxy |
| ☑ | ApplicationResigstry Service Proxy | Application Registry Proxy |
| ☑ | BDC Service Application Proxy | Business Data Connectivity Service Application Proxy |
| ☑ | Excel Service Application | Excel Services Application Web Service Application Proxy |

Custom groups that are created in Central Administration are not reusable across multiple Web applications. Each time that you select **custom** when you create a Web application, you are selecting service applications only for the Web application that you are creating.

## Logical architecture

Service applications are deployed within a single Internet Information Services (IIS) Web site. This is the default behavior and cannot be changed. However, you can customize the configuration of service application groups and the association of Web applications to service application groups.

The following diagram shows the logical architecture for a typical farm deployment.

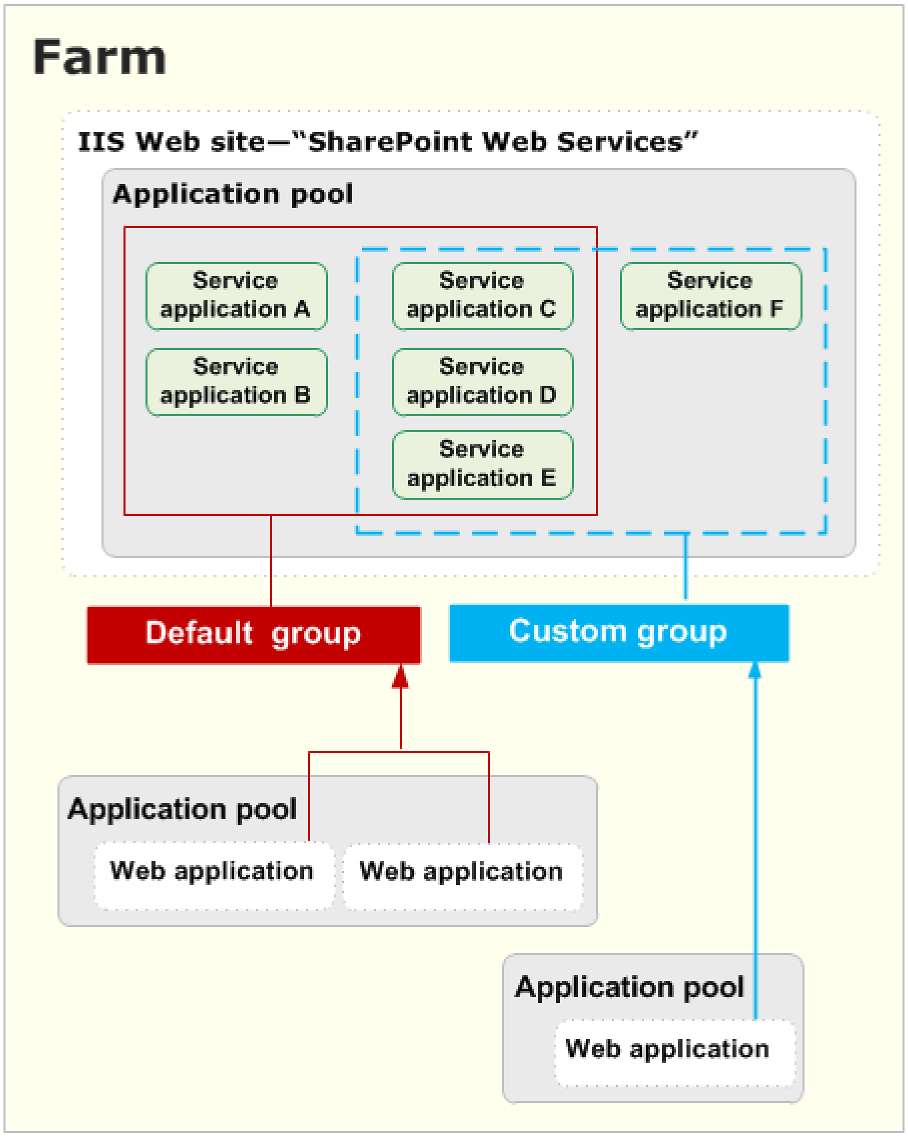

Notice the following characteristics of the farm in the diagram:

- All service applications are contained within the same IIS Web site.
- There are two groups of service applications: the default group and a custom group. Not all service applications have to be included in the default group. In the diagram, Service application F is not included in the default group. It is used only by one Web application.
- Web applications connect either to the default group or to a custom group of service applications. In the diagram, there is one custom group.

Service applications can be deployed to different application pools to achieve process isolation. However, if you want to optimize the performance of your farm, we recommend that you deploy service applications to one application pool.

To achieve physical isolation for a service application, choose or create a different application pool for the service application, as shown in the following diagram.

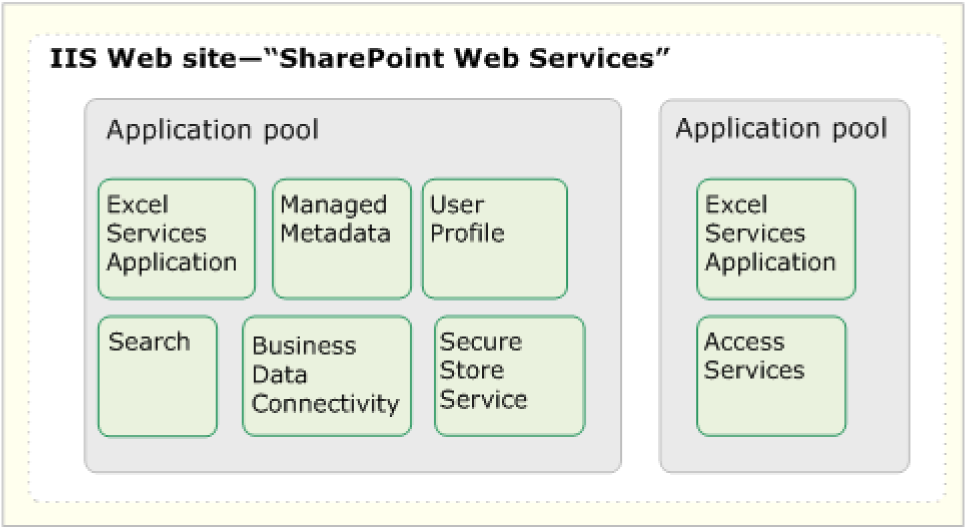
IIS Web site—"SharePoint Web Services"
Application pool
Excel Services Application
Managed Metadata
User Profile
Search
Business Data Connectivity
Secure Store Service
Application pool
Excel Services Application
Access Services

## Connections for service applications

When you create a service application, a connection for the service application is created at the same time. A connection is a virtual entity that connects Web applications to service applications. In Windows PowerShell, these connections are called *proxies*. "Proxy" appears at the end of the type description for connections on the Manage Service Applications page in Central Administration. Some connections might include settings that can be modified. For example, connections for the Managed Metadata service application include several settings, including Term Store Administrators and Default Language.

## Service application administration

Service applications are managed directly in Central Administration rather than through a separate administration site. If needed, service applications can be monitored and managed remotely. Service applications can also be managed and scripted by using Windows PowerShell.

# Deploying service applications across farms

Some service applications can be shared across server farms. Other service applications can be shared only within a single server farm.

The following diagram shows which service applications can be shared across farms and which service applications are limited to a single farm.

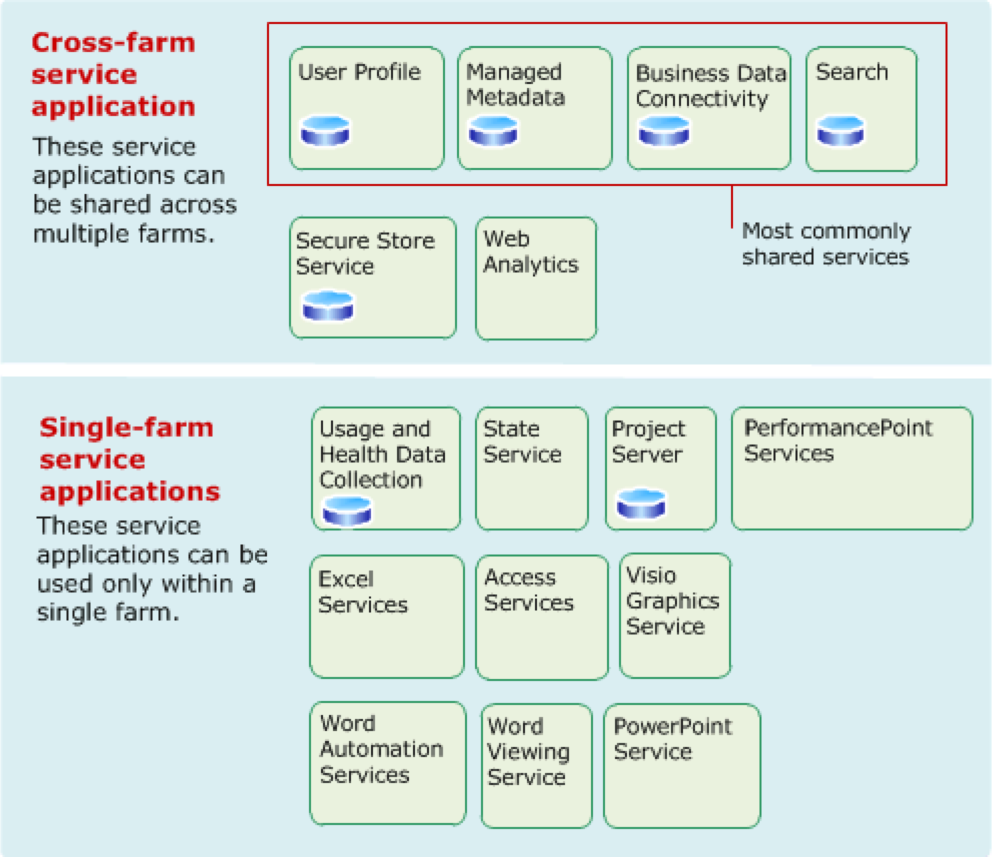
Cross-farm service application
These service applications can be shared across multiple farms.
User Profile
Managed Metadata
Business Data Connectivity
Search
Most commonly shared services
Secure Store Service
Web Analytics
Single-farm service applications
These service applications can be used only within a single farm.
Usage and Health Data Collection
State Service
Project Server
PerformancePoint Services
Excel Services
Access Services
Visio Graphics Service
Word Automation Services
Word Viewing Service
PowerPoint Service

## Design guidance

The following guidance applies to sharing service applications across farms:

- Service applications that support sharing across farms can be run in a central farm and consumed from other farms.
- Each Web application can be configured to use service applications from different farms. For example, you can share the User Profile service across Web applications in several server farms, while at the same time you can configure some service applications, such as the Business Data Connectivity service, to be used locally.
- In large environments, computing-intensive service applications can be run in a central farm to minimize administrative overhead and to scale out easily and efficiently as requirements grow. For more information, see Enterprise services farms, later in this article.

## Deploying cross-farm services

Sharing service applications across farms requires several steps:

1. Configure trusted farms.

   Ensure that farms have exchanged certificates to trust one another. Export the certificate to a file, and back up the file before you connect to cross-farm services.

2. Publish the service applications.

   To share a service application across farms, you first publish the service.

3. Connect to cross-farm service applications.

   To consume a service that is published by a remote farm, create a connection to the service. This process prompts you to enter the URL of a published service, which is displayed during the publish process. A connection on the local farm is created to connect to the service application on the remote farm.

If the server farms are located in two domains, the User Profile service application requires both domains to trust each other. For the Business Data Connectivity and Secure Store service application administration features to work from the consuming farm, the domain of the publishing farm must trust the domain of the consuming farm. Other cross-farm service applications work without a trust requirement between domains.

For more information about how to configure services for use across farms, see Connect to a service application on a remote farm (SharePoint Server 2010).

# Planning considerations for services that access external data sources

Some service applications can access external data sources. Service applications that access external data sources by using a delegated Windows identity put additional requirements on the environment. For these service applications, external data sources must reside within the same domain as the SharePoint Server 2010 farm where the service applications are located or the service application must be configured to use the Secure Store Service.

The following service applications access external data sources by using a delegated Windows identity:

- Excel Services
- PerformancePoint Services
- InfoPath Forms Services
- Visio Services

Service applications that access external data sources by using a delegated Windows identity can be configured to use the Secure Store Service as an alternative. The Secure Store Service stores and maintains user or service credentials. Service applications can use stored credentials to authenticate to a data source directly.

If the Secure Store Service is not used and external data sources do not reside within the same domain, authentication to the external data sources will fail. If farm servers are split between two domains, the application servers must reside in the same domain as the external data sources.

The following service applications and products are not affected by these requirements:

- Business Data Connectivity service and Microsoft Business Connectivity Services
- Access Services
- Microsoft SQL Server PowerPivot for Microsoft SharePoint
- Microsoft SQL Server Reporting Services (SSRS)
- Microsoft Project Server 2010

# Example architectures

The rest of this article provides example architectures for common deployment scenarios.

# Single farm, single service application group

In an architecture that includes a single farm and a single service application group, the default service application group is used for all Web applications in the farm. All sites have access to all of the service applications that are deployed in the farm.

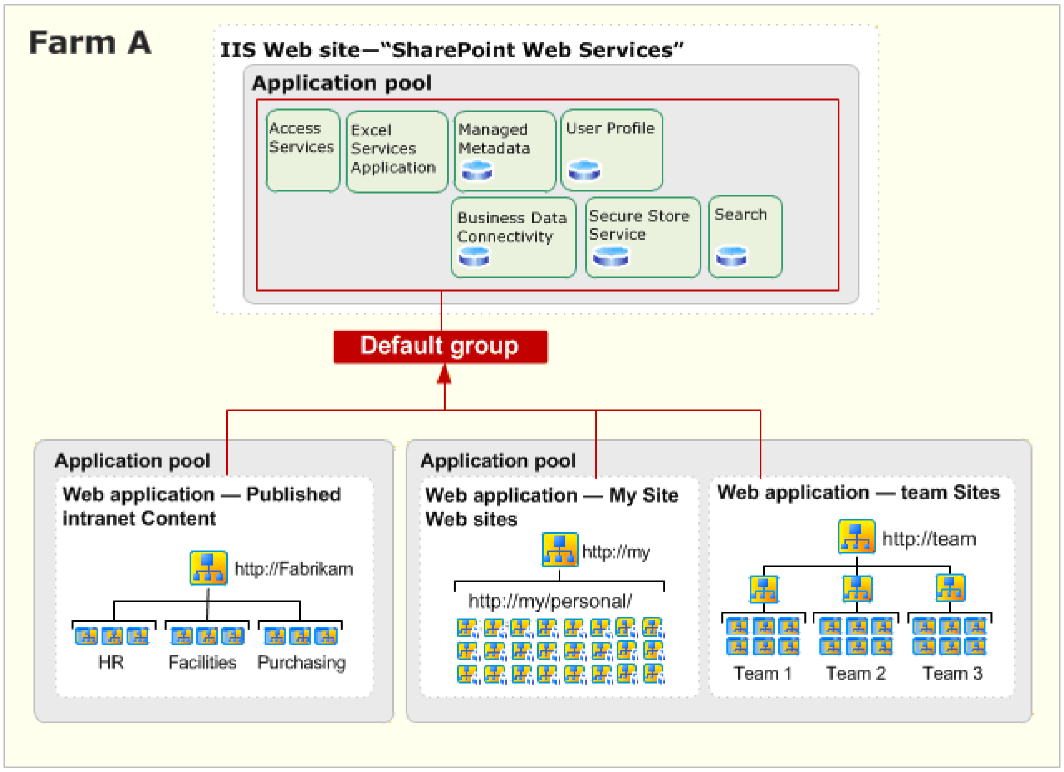

## Advantages

This architecture provides the following advantages:

- It is the simplest architecture to deploy.
- All service applications are available to all Web applications.
- Farm resources are used most efficiently.
- All service applications are managed centrally.

### Disadvantages

There are several tradeoffs to consider with this architecture:

- You cannot isolate service application data.
- Individual departments or teams cannot manage service applications on their own.

### Recommendations

The architecture that includes a single farm and a single service application–group is the recommended configuration for most organizations, at least initially. This configuration works well when you want to host many sites for a single company on the same farm.

Use this configuration to meet the following goals:

- You want to optimize the resources required to run service applications within a farm.
- You are sharing content and profile data across sites that otherwise require process isolation for performance or security reasons.

## Single farm, multiple service application groups

If teams require dedicated service applications, build an architecture by using one or more custom groups of service applications. Follow these guidelines:

- Deploy specific service applications for dedicated use by one or more teams within an organization.
- Ensure that the dedicated service applications are not also included in the default group.
- Create one or more Web applications that use a custom group of service applications. The SharePoint administrator selects the service applications that are included in the custom group.

In the following diagram, Farm B shows an example architecture with two groups of service applications. In this example, the Finance team requires a dedicated Excel Services application. Access Services is also deployed for this team.

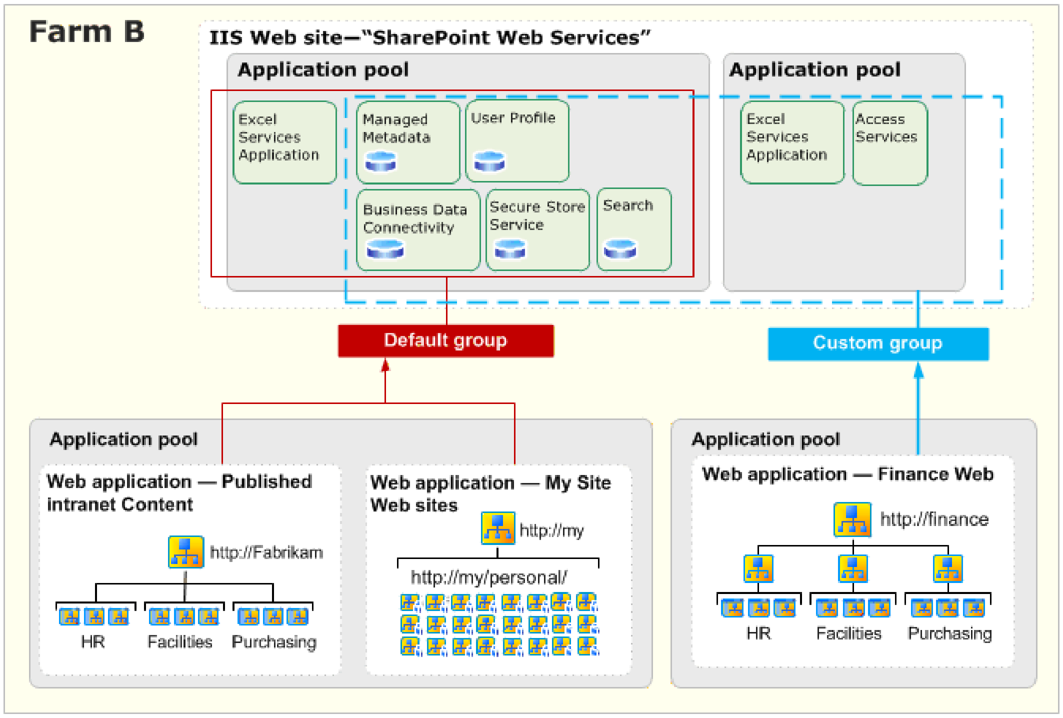

You can create more than one custom service application group. In the following diagram, two custom groups are created in Farm C. Building on the architecture for Farm B, dedicated Managed Metadata and Business Data Connectivity service applications are deployed to the farm for use by the HR department. This results in a second custom service application group, in addition to the first dedicated service application group that was created for the Finance team.

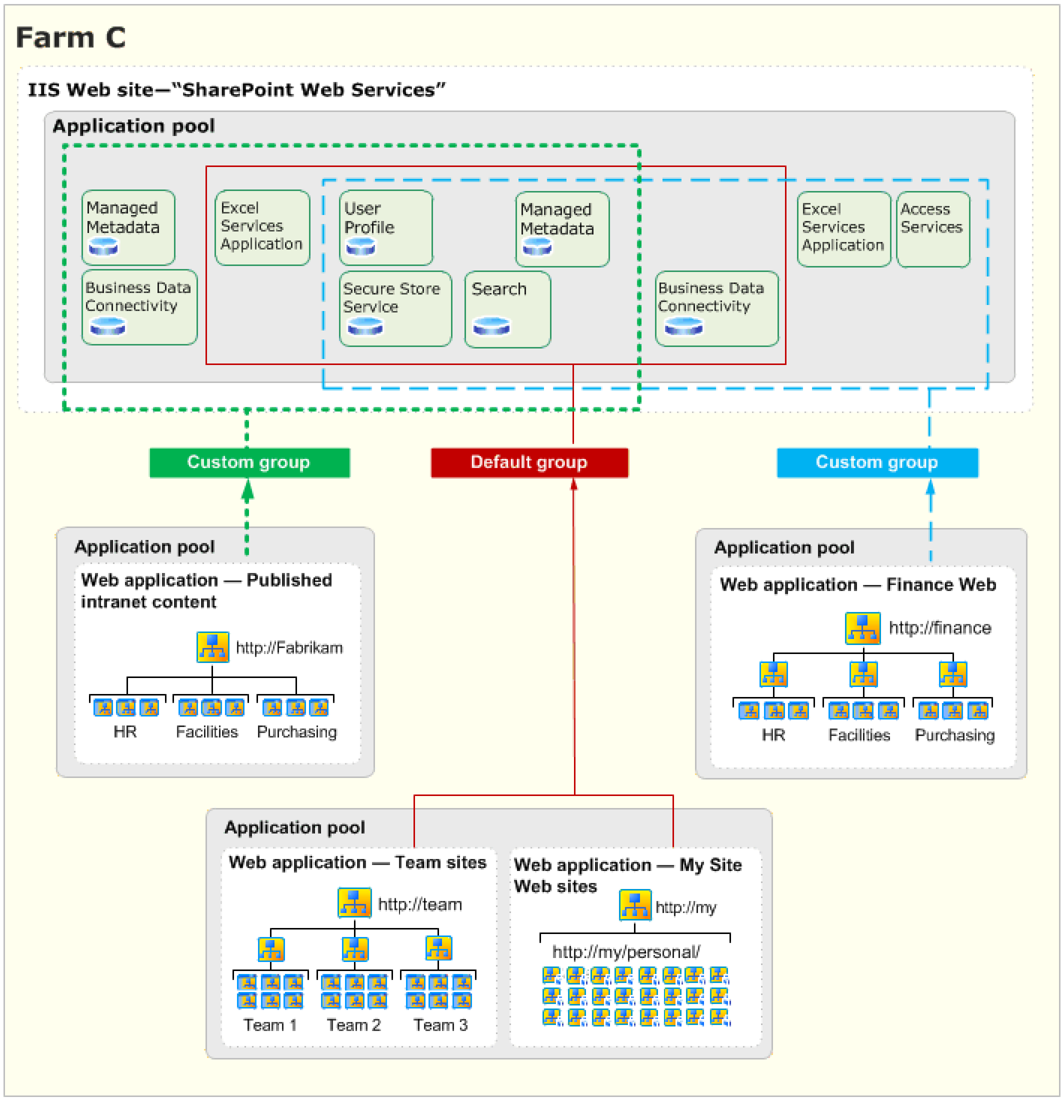

Service applications that are deployed for dedicated use can share the same application pool or be deployed to a separate application pool for additional isolation. The design of Farm B (two diagrams previous) achieves process isolation for service applications that are deployed for the Finance team by putting these service applications in a dedicated application pool. For Farm C, shown in the preceding diagram, one application pool is used for all service applications; in this architecture, service applications are deployed to optimize performance instead.

## Connecting to multiple Managed Metadata service applications

A service application group can include multiple Managed Metadata service applications. For example, in the diagram of Farm C, the custom group that is highlighted in green includes two Managed Metadata service applications.

In this scenario, the sites within the Web applications display taxonomy, social tagging, and other features from both Managed Metadata service applications. Unlike other cross-farm services, Web parts by default include data from multiple Managed Metadata service applications.

For information about how to manage multiple Managed Metadata service applications, see Managed metadata service application overview (SharePoint Server 2010).

## Advantages

Architectures that include multiple service application groups provide the following advantages:

- They accommodate multiple organizational goals in the same farm.
- Service data can be isolated.
- Individual teams or departments can manage the service applications that have been dedicated for their use.
- Sites can be configured to use a subset of service applications.

## Disadvantages

The tradeoffs in architectures that use more than one group of service applications include the following:

- They are more complex to configure and manage.
- Farm resources are consumed to support multiple instances of some service applications, which can affect performance.

## Recommendations

Architectures that include multiple service application groups work well for companies that have divisions or teams that require dedicated service applications or isolated service data, or for sites that are set up with a narrower scope, such as partner collaboration.

Additionally, when multiple groups of service applications are configured, teams and sites can consume services that are offered enterprise-wide, such as profile and search services, while at the same time isolating the use of targeted services for security or performance reasons.

Service applications that are typically deployed for dedicated use by a specific team or department include:

- **Excel Services** To optimize performance for a targeted team or to isolate sensitive data.
- **Managed Metadata** To allow a team or department to manage their own taxonomy, hierarchies, keywords, and so on. SharePoint Server 2010 combines the results of multiple Managed Metadata service applications, so taxonomies, content types, and other elements can be shared across an organization.
- **Business Data Connectivity** Individual teams or departments can integrate with their own line-of-business data systems and keep the data isolated from the rest of the organization.

In some cases, a dedicated group of service applications is configured to narrow the list of services that are used by a Web application. For example, a partner collaboration site can be configured to consume a subset of the service applications that are offered by the farm.

## Enterprise services farms

An enterprise services farm is a server farm that is dedicated to hosting service applications for an organization. The following diagram shows an enterprise services farm that hosts the most frequently deployed cross-farm service applications. It also shows several common kinds of farms that consume services from an enterprise services farm.

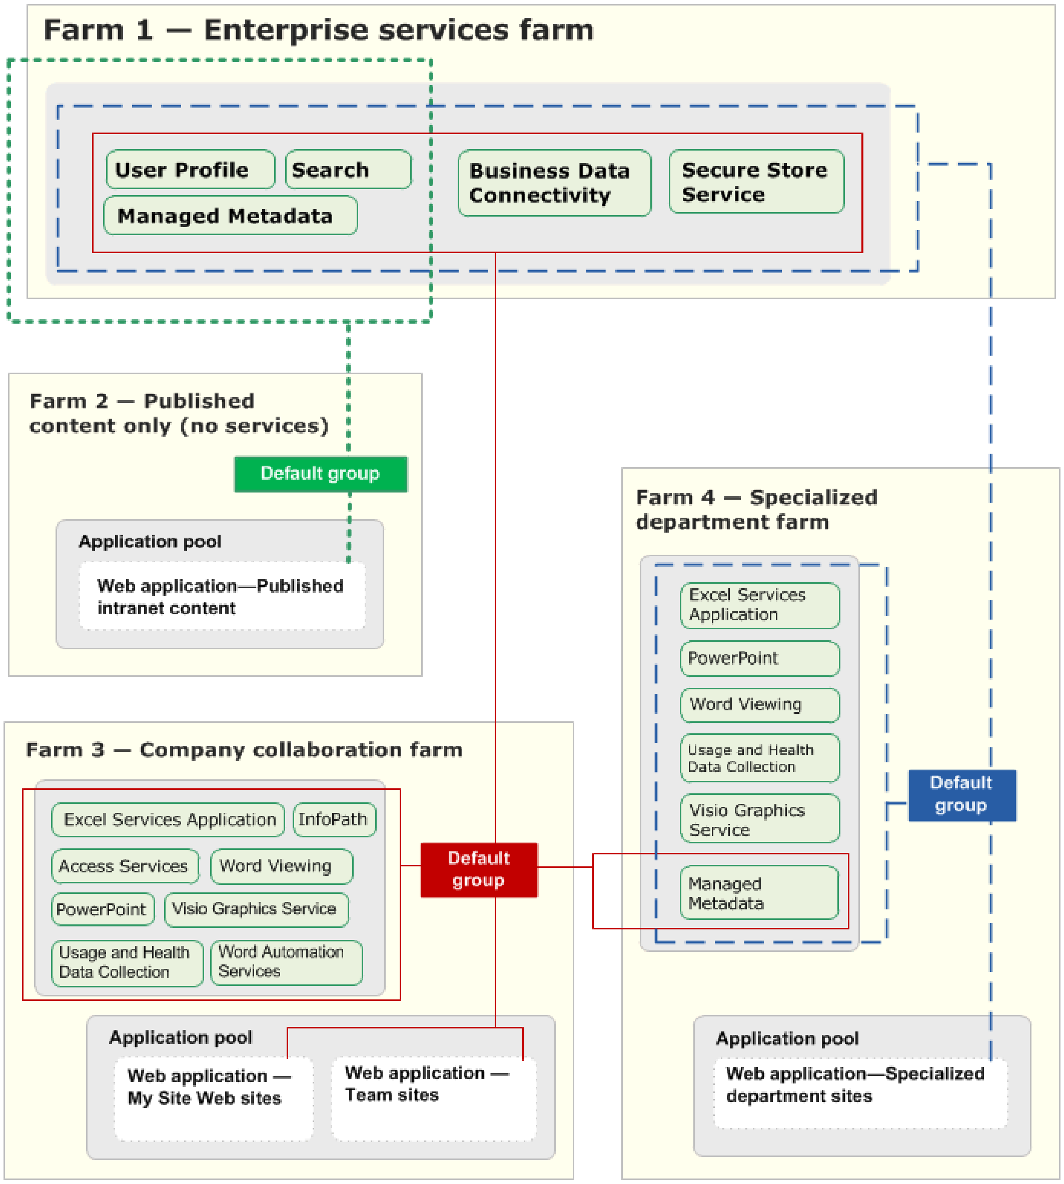

Farm 1 — Enterprise services farm
User Profile
Search
Managed Metadata
Business Data Connectivity
Secure Store Service
Farm 2 — Published content only (no services)
Default group
Application pool
Web application—Published intranet content
Farm 3 — Company collaboration farm
Excel Services Application
InfoPath
Access Services
Word Viewing
PowerPoint
Visio Graphics Service
Usage and Health Data Collection
Word Automation Services
Default group
Application pool
Web application — My Site Web sites
Web application — Team sites
Farm 4 — Specialized department farm
Excel Services Application
PowerPoint
Word Viewing
Usage and Health Data Collection
Visio Graphics Service
Managed Metadata
Default group
Application pool
Web application—Specialized department sites

The rest of this section describes the other farms in the diagram. Farm 2, Farm 3, and Farm 4 represent the kinds of farms that are most likely to consume services from an enterprise services farm.

## Published content–only farms (all service applications are remote)

You can deploy a server farm without deploying any service applications locally. In Farm 2, no service applications are hosted locally. All service applications are consumed from a separate farm.

This configuration works well for content that is published. It reduces the administrative efforts required to host a published content farm and allows an organization to take advantage of centrally managed service applications.

Use this configuration when you have the following goals:

- You want to optimize the resources within a farm for hosting content, instead of running service applications.
- You are integrating with organization-wide profiles, metadata, search, and other centrally managed resources.

## Collaboration farms (mix of local and remote service applications)

Farm 3 represents a farm that is optimized for collaboration. All service applications that cannot be shared across farms are hosted locally. These include the client-related service applications, which are important for collaboration. Cross-farm service applications are consumed from an enterprise services farm (Farm 1).

Farms can consume services from more than one remote farm. In the diagram, Farm 3 also consumes the Managed Metadata service from a specialized department farm (Farm 4) to integrate with this department's autonomously managed taxonomy, social tagging, and other features.

If there are multiple Managed Metadata service applications, one of the service applications must be designated as the primary service application that hosts the corporate taxonomy. All other instances of this service application are then secondary, and they provide additional data to the primary service application data. Unlike other cross-farm services, Web parts by default include data from multiple Managed Metadata service applications.

This is the recommended configuration for companies that host multiple farms to meet business needs. Use this configuration to meet the following goals:

- To optimize administrative and farm resources at the enterprise level for hosting services (Farm 1).
- To optimize resources at the farm level for hosting collaboration sites (Farm 3).
- To integrate with organization-wide profiles, metadata, search, and other centrally managed resources.
- To integrate with metadata produced by a specialized team or department (Farm 4).

### Farms for specialized departments (mix of local and remote service applications)

Some teams within an organization might require a separate deployment of specific services for the following reasons:

- To ensure data isolation (such as Business Data Connectivity data).
- To provide the ability to autonomously manage service applications (such as Managed Metadata).

Farm 4 provides an example. The characteristics of this farm include the following:

- It consumes centrally managed service applications that include Managed Metadata.
- It also includes its own Managed Metadata service application, so this team can autonomously manage its own metadata. Because this service application is shared, the metadata from the rest of the organization can be integrated with this metadata.

Use this configuration to meet the following goals:

- Allow a specialized team or department to manage metadata on their own.
- Ensure that specific service data is isolated and managed separately from the rest of the organization.

## Specialized service farms

Consider deploying specialized service farms to optimize farm resources for specific service applications. This allows you to scale out the server farm and to scale up the hardware to optimize performance for a specific service application.

The primary service application that might warrant a dedicated services farm is Search. Search has unique performance and capacity requirements. By offloading the Search service application to a dedicated farm, resources can be optimized for the remaining cross-farm service applications.

The following diagram shows two centralized services farms. One farm is optimized for Search. The other farm hosts all other cross-farm service applications.

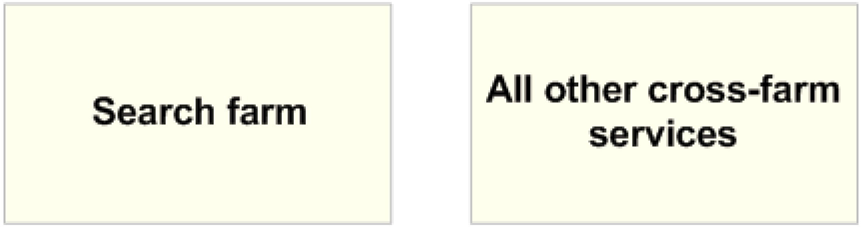

# Cross-organization farms

Service applications can be shared across any farm, not only enterprise services farms. Consider sharing service applications across farms in the following scenarios.

Scenario A: To provide enterprise-wide service applications without a dedicated enterprise services farm, as shown in the following diagram.

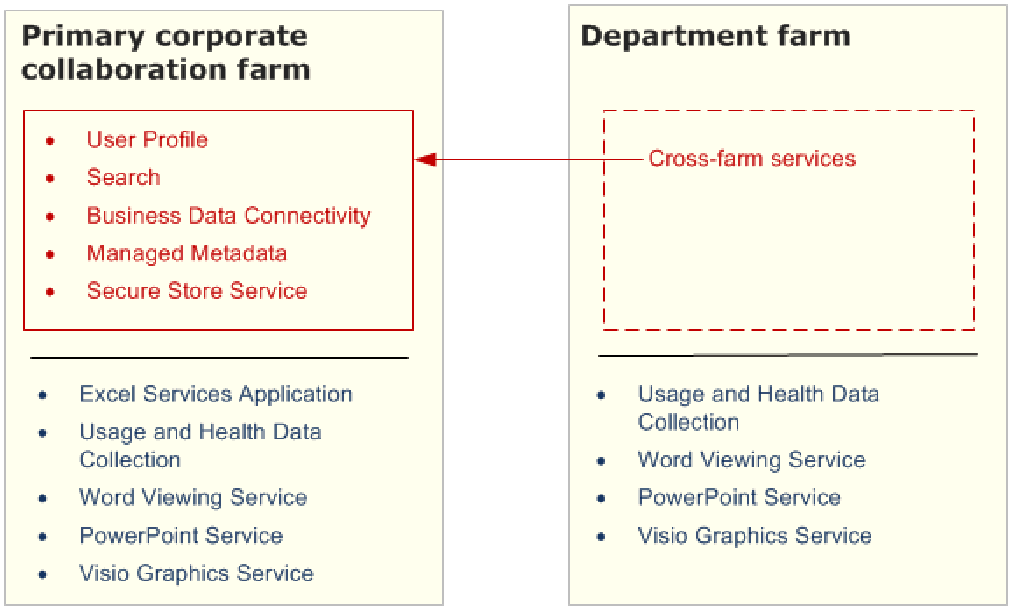

Scenario B: To share resources across farms and to avoid deploying redundant service applications, as shown in the following diagram.

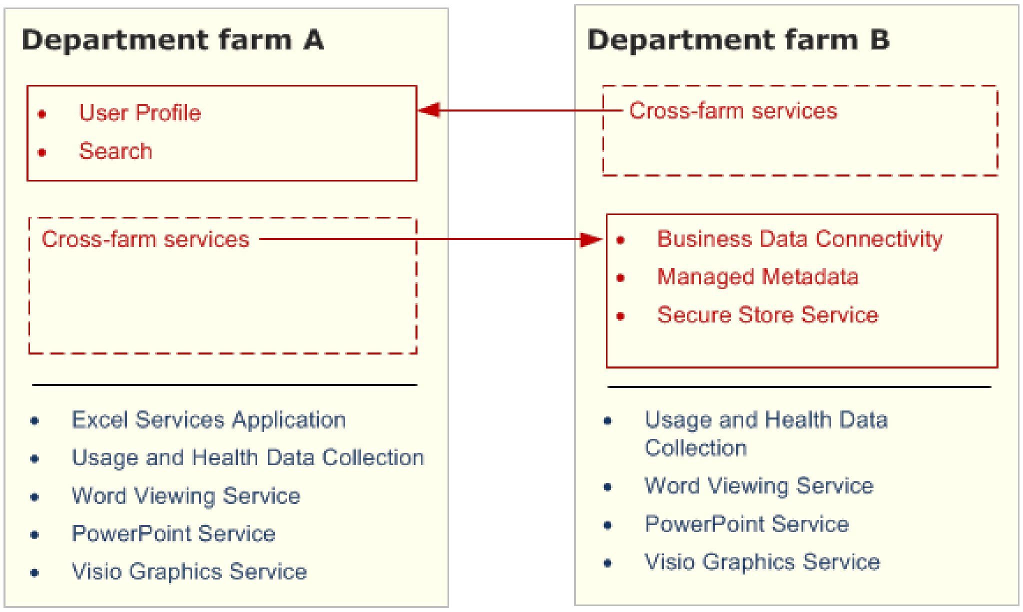
Department farm A
User Profile
Search
Cross-farm services
Excel Services Application
Usage and Health Data Collection
Word Viewing Service
PowerPoint Service
Visio Graphics Service
Department farm B
Cross-farm services
Business Data Connectivity
Managed Metadata
Secure Store Service
Usage and Health Data Collection
Word Viewing Service
PowerPoint Service
Visio Graphics Service

# Logical architecture components (SharePoint Server 2010)

There are a variety of ways you can arrange the components in a logical architecture design. Each of the components presents different opportunities for sharing and isolation. Before you begin your logical architecture design:

- Know what your sharing and isolation goals are.
- Evaluate the tradeoffs for each choice.

Each section in this article describes a particular logical architecture component and discusses the following considerations for that component: capacity, sharing and isolation, configurable items, administration, and planning recommendations.

In this article:

- Server farms
- Service applications
- Application pools
- Web applications
- Zones
- Policy for a Web application
- Content databases
- Site collections
- Sites
- Host-named site collections
- My Sites

## Server farms

A server farm represents the top-level element of a design. Individual server farms provide physical isolation.

Several criteria that are determined by your organization might affect the number of server farms that are required, including:

- Heavy use of services might warrant one or more dedicated services farms.
- Separate operational divisions of responsibility.
- Dedicated funding sources.
- Separate datacenter locations.
- Industry requirements for physical isolation between sites.

However, you can satisfy many isolation requirements on a single server farm. For example, you can use different Internet Information Services (IIS) application pools with different process identities to achieve isolation at the process level for both sites and service applications.

In addition to isolation requirements that might require more than one server farm, an organization might implement multiple server farms to satisfy performance and scale goals, licensing requirements, or a publishing environment.

# Service applications

A service application provides a resource that can be shared across sites within a farm or, in some cases, across multiple farms.

In SharePoint Server 2010, services are no longer contained within a Shared Services Provider (SSP). Instead, the infrastructure for hosting services moves into Microsoft SharePoint Foundation 2010 and the configuration of service offerings is much more flexible. Individual services can be configured independently and third-parties can add services to the platform.

You can deploy only the services that are needed to a farm. Services that are deployed are called service applications.

Service applications are associated with Web applications. Each service application can be configured differently:

- Web applications can be configured to use only the services that are needed, rather than the entire set of services that are deployed.
- You can deploy multiple instances of the same service in a farm and assign unique names to the resulting service applications.
- You can share service applications across multiple Web applications within the same farm.
- Some service applications can be shared across farms.

## Capacity

There is no recommended limit for the number of service applications in a single farm.

## Sharing and isolation

Service applications can be shared in two ways:

- Sharing the service application and the service data. This is the default behavior for services that are shared across Web applications. For example, search results are shared across Web applications that consume the same search application.
- Sharing only the service application, but isolating the data by deploying the service in partitioned mode. In a hosted environment, you can deploy service applications in partitioned mode by using Windows PowerShell. Each tenant’s data is stored in a separate partition in the database for the service. A tenant’s subscription ID is used to map the tenant’s service data to their sites. For

example, if you deploy the search service in partitioned mode, each tenant will only see search results for their own content.

**Note:**
Not all service applications support partitioning.

Conversely, service applications can be isolated in two ways:

- Deploying multiple service applications in separate application pools to achieve process isolation of services and service data. For example, a finance team might warrant a separate and dedicated Business Data Connectivity application.
- Deploying services in partitioned mode. This approach works well in hosted environments in which tenants will never share service data. However, it might not be practical in environments where there is a mixture of needs for shared and isolated service data.

If needed, you can additionally isolate service applications by deploying them to separate application pools to achieve process isolation. However, application pools are a limited resource and farm performance is affected if too many application pools are used. For more information, see Application pools in this article.

## Configurable items

The following table displays configurable items that contribute to isolation and sharing.

| Item | Description |
|---|---|
| Default group | By default, all service applications are included in the default group. You can add and remove service applications from the default group at any time, including when you create one.<br>When you create a Web application, you can select the default group or you can create a custom group of services. You create a custom group of services by selecting only the service applications that you want the Web application to use. |
| Connection (proxy) | When you create a service application, a connection for the service application is created at the same time. A connection is a virtual entity that connects Web applications to service applications. Some service applications, such as the Managed Metadata Service, store settings in the connections. In Windows PowerShell, connections are called proxies. |

| | |
|---|---|
| Service application permissions | You can delegate management of service applications to other users by granting them permissions to one or more of the service applications. |
| Trusted My Site host locations | In organizations where multiple User Profile Service applications are deployed, this feature ensures that users create a My Site in the location that is intended for their profile. This feature prevents users from creating multiple My Sites across an organization. |

## Administration

Configuration and management of service applications can be delegated to administrators who specialize in managing individual services, such as search, user profiles, and managed metadata.

In a hosted environment, tenants can manage some of the service settings for their organization.

## Planning recommendations

Configure service applications either to share resources across multiple Web applications or isolate content.

For example, multiple sites that reside in different Web application and application pools can be unified by sharing services in the default proxy group to provide taxonomy, content type, and profile sharing across an intranet. This provides for personalization and enterprise-wide standardization across many sites and applications. This choice provides an example of balancing process isolation (by implementing separate Web applications and application pools) with the business need to share information and leverage profile data across the applications.

You can also configure service applications to enhance your overall isolation goals. For example, using a dedicated set of services for partner collaboration ensures that partner users cannot search on or access sensitive information within your intranet environment. You can configure individual services to further isolate content between site collections. For example, you can:

- Limit search scopes to the individual site collections.
- Configure the User Profile service to only display users that are part of the same organizational unit in Active Directory Domain Services (AD DS).
- Use the Stsadm command-line tool to configure the People Picker to display only users that are members of the site collection.

When you design your services strategy for an organization, consider the ways in which you can configure the individual services to enhance your overall content sharing or isolation goals.

When you design a services strategy for a hosting environment, determine which services will be available and partitioned.

# Application pools

In Internet Information Services (IIS) 7.0, an application pool is a group of one or more URLs that are served by a worker process or set of worker processes.

When you create Web applications and services in SharePoint 2010 Products, you select an application pool to use or you can create a new application pool. Each application pool has its own worker process and can have a separate identity (security account) which prevents two processes from interacting.

## Capacity

The memory overhead of an application pool is 30-50 megabytes (MB) plus any memory for the applications running in the application pool process space. The various application demands usually quickly drives the memory usage of an application pool to 800 MB or larger. The limit for the number of application pools is influenced by the available memory on the system. That is, the number of application pools is dictated by the following two factors:

- Available addressable memory.
- The amount of memory consumed by applications running in the application pool.

The general guideline for acceptable performance is to use eight or fewer application pools.

## Sharing and isolation

IIS 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 can help to prevent an exploit on one site that enables the attacker to inject malicious code that can attack sites in different application pools. More importantly, this strategy isolates code that introduces memory issues or other issues so that the problematic code does not affect all applications.

## Configurable items

Using a separate application pool identity for each application pool is recommended, if needed, for security and reasons of isolation.

## Administration

If separate identities are used for each application pool, each identity will have to be maintained.

## Planning recommendations

Practically speaking, consider using a dedicated application pool for each of the following reasons:

- To separate authenticated content from content that is primarily anonymous.

- To isolate applications that store passwords for and interact with external business applications, for example, Business Data Connectivity connections.

# Web applications

A Web application is an IIS Web site that is created and used by SharePoint 2010 Products. A Web application can be extended up to four times to create four additional zones in SharePoint 2010 Products, resulting in up to five IIS Web sites that are associated with a single Web application, each IIS Web site associated with a different zone. You can assign each Web application a unique domain name. For more information, see Zones in this article.

## Sharing and isolation

Each Web application has a unique domain name, which helps to prevent cross-site scripting attacks.

## Configurable items

The following table displays configurable items that contribute to isolation and sharing.

| Item | Description |
|---|---|
| Service applications | Service applications are associated with Web applications. When you create a Web application, you can select the default proxy group (default set of service applications), or you can specify a custom set of service applications for the Web application. All sites within a Web application consume services from the same service applications. A service application can provide services for more than one Web application, thus sharing content and profile data across the Web applications. |
| Zones | Within a Web application, you can create up to five zones. Use zones to enforce different access and policy conditions for large groups of users. |
| Policy for Web application | Create a policy to enforce permissions across one or more zones in a Web application. A policy can be created for a specific user or user group. For more information, see Policy for a Web application in this article. |

## Administration

Ongoing administration of Web applications is not significant.

## Planning recommendations

Generally speaking, use dedicated Web applications to:

- Separate content available to anonymous users from content available to authenticated users.
- Isolate users. For example, you can ensure that partners do not have access to intranet content by placing partner sites in a separate Web application.
- Enforce permissions through the use of policies. For example, you can create a policy to explicitly deny write access to one or more groups of users. Policies for a Web application are enforced regardless of permissions configured on individual sites or documents within the Web application.
- Optimize database performance. Applications achieve better performance if they are placed in content databases with other applications with similar data characteristics. For example, the data characteristics of My Sites typically 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.
- Optimize manageability. Because creating separate Web applications results in separate sites and databases, you can implement different limits for each site's Recycle Bin, expiration, and size, and negotiate different service-level agreements. For example, you might allow more time to restore sites that are not critical to your business.

# Zones

Zones represent different logical paths (URLs) of gaining access to the same Web application. Within each Web application, you can create up to five zones using one of the available zone names: Default, Intranet, Internet, Custom, or Extranet. Each name can only be selected once per Web application. Each zone is represented by a different Web site in IIS.

The Default zone is the zone that is first created when a Web application is created. The other zones are created by extending a Web application.

## Capacity

You can create up to five zones within a Web application. Typically, zones are coordinated across Web applications so that zones of the same name are configured for the same users.

## Sharing and isolation

Zones provide a method of partitioning users by:

- **Authentication type**: Each zone can be configured to use a different authentication provider, enabling you to share the same content across partner companies.
- **Network zone**: Each zone can be configured to accommodate users entering from a different network zone, such as an extranet or the Internet.
- **Policy permissions**: You can explicitly allow or deny read or write access to content per zone based on a user account or a group account.

## Configurable items

The following table displays configurable items that contribute to isolation and sharing.

| Item | Description |
|---|---|
| Authentication provider | Each zone can be configured to use a different authentication provider. |
| Anonymous access | Turn anonymous access on or off per zone. |
| Secure Sockets Layer (SSL) | Turn SSL on or off per zone. |
| Public URL and alternate access mapping | Specify the domain name users will type to access content in the Web application. Alternatively, use alternate access mapping to map user-friendly or zone-appropriate URLs to the default URL (server name and port) for each zone. Alternate access mapping provides support for off-box termination of SSL. Off-box termination of SSL is when a proxy server terminates an SSL request and then forwards the request to a Web server by using HTTP. In this case, alternate access mappings can be configured to return these requests using SSL, thus maintaining secure communication between the client and the proxy server. |
| Policy for Web application | Create a unique set of policies for each zone within the Web application. If you have a special group of users that require exceptions to your overall security policy, consider using a separate zone to accommodate these users. |

## Administration

If you use alternate access mapping, consider that all public URLs require Domain Name System (DNS) entries to map the public URLs to the IP address of the load balancer used for the farm.

## Planning recommendations

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 some of the planning recommendations and requirements for zones, including the default zone.

- Administrative e-mail is sent with links from the Default zone. This includes e-mail to owners of sites that are approaching quota limits. Consequently, users who receive administrative e-mail messages and alerts must be able to access links through the Default zone. This is especially important for site owners.
- Host-named site collections are only available through the Default zone. All users who are intended to access host-named site collections must have access through the Default zone.
- The Default zone must be the most secure zone. This is because when a user request cannot be associated with a zone, the authentication and policies of the Default zone are applied.

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

- User requests can be initiated from several different networks, such as the internal network, a partner company network, or the Internet.
- Users consume content across multiple Web applications. For example, an intranet environment might include sites that are hosted in several different Web applications. Additionally, employees might have access to both the intranet content and to partner collaboration content.

In an extranet environment, ensure that the following design principles are followed:

- Configure zones across multiple Web applications to mirror each other. The configuration of authentication and the intended users should be the same. However, the policies associated with zones can differ across Web applications. For example, ensure that the Intranet zone is used 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.
- Configure alternate access mappings appropriately and accurately for each zone and each resource.

# Policy for a Web application

A policy for a Web application enforces permissions on all content within a Web application, enabling you to set security policy for users at the Web application level. The permissions in a policy override all other security settings that are configured for sites and content.

You can configure policy based on users or user groups in AD DS, but not SharePoint groups. A policy can be defined for the Web application in general or just for a specific zone.

## Capacity

There are no capacity restrictions that apply to policies for Web applications.

## Sharing and isolation

A policy for a Web application provides a method of setting permissions based on users and the zone that they access content through.

For example, by using a policy, you can:

- Allow Help desk staff access to all content.
- Deny write access to partners or vendors.
- Deny access to secure data to a group of users regardless of how site owners configure permissions.
- Ensure that the crawl account has access to crawl all content.

## Configurable items

The following table displays configurable items that contribute to isolation and sharing.

| Item | Description |
|---|---|
| | |
| User policy | Create a policy that applies to users or user groups:<br>• The policy can be applied to all zones or one zone.<br>• You can enter user names, group names, or e-mail addresses.<br>• Specify the permissions that you want to apply to the policy.<br>You can modify the default permission levels or create new permission levels by clicking Permission policy when you create the policy in Central Administration. |
| Anonymous policy | If anonymous access is enabled for the Web application or one or more zones, then you can create a policy that applies to all anonymous |

| | |
|---|---|
| | users. The default policy settings are:<br>• None: No policy<br>• Deny write: No write access<br>• Deny all: No access<br>Anonymous user policy levels cannot be modified. |
| Permission policy | Edit the specific permissions associated with one of the default permission levels, or create a new permission policy level. Additionally, you can specify the particular permissions that allowed or denied for site collections and sites.<br>After creating a new permission policy level, you can create a user policy that uses the permission policy. |

## Administration

Ongoing administration of policies for Web applications is not significant.

## Planning recommendations

Because policies are managed centrally, consider using policies to manage large groups of users, rather than individual users.

# Content databases

By default, all content for a Web application is stored in one content database. You can separate content into multiple content databases at the site collection level. A content database can include one or more site collections. A single site collection cannot span multiple databases. Backing up and restoring sites takes place at the content database level.

## Capacity

The guideline for acceptable performance is to implement 100 or fewer content databases per Web application.

## Sharing and isolation

Planning for databases enables you to either optimize for efficiency (multiple site collections sharing a database) or isolation (one database per site collection).

Achieve scale efficiency by managing databases to the maximum target size. In this case, you configure database settings to add new site collections to existing databases until the maximum number of site collections has been reached. You calculate the maximum number of site collections by estimating the average or maximum size of site collections divided into the maximum target size for the database. This approach works well when you expect a large number of small site collections, such as My Sites.

Achieve isolation of content between teams or projects by limiting a database to one site collection. This approach enables you to independently manage the content of individual teams. For example, you can independently manage each team's database for backup, recovery, and migration. This approach provides the opportunity to implement different service-level agreements for different teams or projects. This approach also enables you to manage content to the lifecycle of a project. For example, you can archive a database when a project is completed.

## Configurable items

The following table displays configurable items that contribute to isolation and sharing.

| Item | Description |
|---|---|
| Database server | Specify which SQL Server computer a content database is created on. |
| Failover server | You can choose to associate a content database with a specific failover server that is used in conjunction with SQL Server database mirroring. |
| Capacity settings | You can specify the number of sites that can be created before a warning event is generated and the maximum number of sites that can be created in each database. |

## Administration

A manageable database administration plan balances the number of databases with the resources required to manage the databases.

Administration of databases includes:

- Creating new databases for new team sites or site collections that require dedicated databases.
- Monitoring database sizes and creating new databases when target sizes are approached.
- Backing up and restoring databases.

### Planning recommendations

Choose one of the following two approaches:

- Establish target sizes for content databases with appropriate size-warning thresholds. Create new databases when size-warning thresholds are reached. With this approach, site collections are automatically added to the available database or databases, based on target sizes alone.
- Associate site collections with specific content databases. This approach enables you to place one or more site collections in a dedicated database that can be managed independently from other databases.

If you want 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 new database.
- Apply the following database capacity settings on the Manage Content Database Settings page on the SharePoint Central Administration Web site:
  - Number of sites before a warning event is generated = 0
  - Maximum number of sites that can be created in this database = 1
- Add a group of site collections to a dedicated database by performing the following steps:
  a. Add a content database for the Web application and ensure that the database status is set to Ready.
  b. 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.
  c. Create the site collections. They are automatically added to the online database.
  d. Set the status of all other databases back to Ready.

## Site collections

A site collection is a set of Web sites that have the same owner and share administration settings. Each site collection contains a top-level Web site and can contain one or more subsites.

### Capacity

The recommended guideline for acceptable performance is to implement fewer than 50,000 site collections per content database; however, performance can be affected at about 10,000 site collections. Scaling out by distributing site collections across multiple database servers provides additional storage capacity and throughput.

### Sharing and isolation

Site collections introduce several sharing and isolation opportunities that affect permissions, navigation, and feature deployment.

The following items can be shared within a site collection and cannot be shared across site collections (except items that are stored in a file system, such as features in the _layouts directory):

- Master pages
- Page layouts
- Images
- Site templates

Additionally, permissions and navigation are isolated at the site collection level in the following ways:

- Subsites within a site collection can inherit permissions from the top-level site.
- Site collections cannot inherit permissions from other site collections.
- There is no built-in navigation from one site collection to another.

Finally, SharePoint Server 2010 aggregates search results across site collections based on a user's permissions, regardless of the number of site collections or databases (depending on search scopes).

It is important to note that although permissions are enforced on individual sites, the sites are still vulnerable to cross-site scripting attacks from other sites within the domain.

## Configurable items

The following table displays configurable items that contribute to isolation and sharing.

| Item | Description |
|---|---|
| Site collection administrator | You can specify one user to be the primary site collection administrator and one user to be the secondary site collection administrator. In Central Administration, you cannot enter more than one account for these roles, nor can you enter a group account for these roles. |
| Site template | A site template determines which lists and features will be available on your new site. Many aspects of a site can be customized after creation. However, the site template cannot be changed once the site is created. |
| Quota template | You can apply a quota template to limit resources used for a site collection. The following templates are provided:<br>• Personal Site (100 MB)<br>• Team Sites (2,000 MB) |

The following table displays configurable items within a site collection that contribute to isolation and sharing. These settings are available after you create the site collection using the settings in the previous table.

| Item | Description |
|---|---|
| Site collection administrators | You can specify multiple user accounts to be site collection administrators. You cannot add group accounts. |
| Permission level | Add user and group accounts to site collections and specify permission levels for each. |

## Administration

Site collection creation does not require DNS entries (unless you are creating host-named site collections) and can be easily automated or delegated to users. You can create site collections for your team sites centrally, or you can allow users to create their own site collections by using Self-Service Site Management.

Using a dedicated database for a site collection provides the ability to perform backup and recovery at the site collection level.

## Planning recommendations

Site collections bridge logical architecture and information architecture. When you design your site collections, consider the following two design tasks:

- Design consistent URLs across your organization.
- Create logical divisions of content.

Unless you are using host-named site collections, each Web application must have a single root-level site collection. This provides a single URL path into the sites that reside in the Web application. This is also a requirement if you are implementing multiple zones within a Web application. For more information, see Host-named site collections in this article.

Many organizations plan to implement multiple site collections within a Web application for use by different teams or divisions within the organization. Common design goals include the following:

- Maintain a separate and independent site collection for each team.
- Create a unique URL for each team.
- Isolate content between teams.

To satisfy these goals, you can use managed paths to incorporate multiple top-level site collections within a Web application. By defining managed paths, you can specify which paths in the URL namespace of a Web application 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 created 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 associate each of these site collections with a different content database if you want to manage growth and to provide the opportunity to back up and restore these sites separately. An example URL for a site collection created by using this method is http://fabrikam/hr. The limit on site collections created with an explicit inclusion is approximately 100 site collections within a Web application; however, 20 is a good operational maximum. If your organization requires a greater number of site collections, use a wildcard inclusion instead.
- **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 to support Self-Service Site Management, such as My Sites or sites created for partner collaboration. Example URLs for site collections created by using this method are http://partnerweb/sites/project1 and http://partnerweb/sites/project2. In these examples, "http://partnerweb" represents the root-level site collection and "/sites" represents the wildcard inclusion.

# Sites

A site consists one or more related Web pages and other items (such as lists, libraries, and documents) that are hosted inside a site collection.

## Capacity

The guideline for acceptable performance is to implement fewer than 250,000 sites per site collection. You can create a very large total number of Web sites by nesting the subsites. However, a large number of nested subsites can greatly affect the time it takes to upgrade sites. 5,000 sites within a site collection is a good operational target.

## Sharing and isolation

Sites include built-in navigation from one subsite to another within a site collection. There is no built-in navigation from one site collection to another.

As with site collections, separate sites are vulnerable to cross-site scripting attacks from other sites within the domain.

## Configurable elements

From within each site, you can add user or group accounts to the Owners group for that site.

## Administration

You can use a variety of tools to back up and restore individual sites.

# Host-named site collections

Host-named site collections are an option if you want to create multiple root-level site collections within a Web application. For example, administrators for hosting organizations use host-named site collections to create multiple domain-named sites.

There is no special mode, such as host header mode, that is required to create host-named site collections. You create host-named site collections by using Windows PowerShell. Additionally, by using Windows PowerShell, you can use managed paths with host-named site collections (**New-SPManagedPath –HostHeader**).

Host-named site collections give you more control over URLs. However, host-named site collections are only available through the Default zone. User accounts that are configured to authenticate through other zones cannot access host-named site collections.

In SharePoint 2010 Products, host-named site collections support off-box SSL termination. However, only the protocol scheme can be changed off-box (http:// or https://). The reverse proxy server cannot change the host name or the port number (except to switch from the default SSL port to the default HTTP port).

## Capacity

You can create up to 100,000 host-named site collections within a single IIS Web site.

## Sharing and isolation

The independent domain names that result from host-named site collections help prevent cross-site scripting attacks between two sites.

## Administration

Administrative tasks for host-named site collections include the following:

- Add host-named site collections by using Windows PowerShell.
- Each host-named site collection requires a separate DNS entry.

# My Sites

My Sites are special SharePoint sites that are personalized for each user. My Sites are enabled by default as part of the User Profile service, and every user in an organization can create a unique My Site. For information about capacity, sharing and isolation, and administration, see Sites earlier in this article.

# Design sample: Corporate deployment (SharePoint Server 2010)

This article describes a practical implementation of logical architecture components to achieve a workable design. This article is intended to be used together with the following design samples:

- Design sample: Corporate portal with classic authentication
- Design sample: Corporate portal with claims-based authentication

To download either of these models, see SharePoint Server 2010 design samples: Corporate portal with classic authentication or with claims-based authentication (http://go.microsoft.com/fwlink/?LinkId=196872).

In this article:

- About the design samples
- Overall design goals
- Server farms
- Users, zones, and authentication
- Services
- Authoring and publishing alternatives
- Administration sites
- Application pools
- Web applications
- Site collections
- Content databases
- Zones and URLs
- Zone policies

The design samples illustrate a generic corporate deployment of Microsoft SharePoint Server 2010. The design samples apply nearly all of the logical architecture components and illustrate how these are incorporated into the overall design. The two samples illustrate the same services and sites, but incorporate different authentication methods, as follows:

- Classic authentication: This design sample represents a path for upgrading sites from Microsoft SharePoint Server 2010 2007 to Microsoft SharePoint Server 2010. This sample incorporates classic authentication, in which Windows authentication methods are used to access sites. A different zone for each authentication method is used. While Windows authentication is used for SharePoint sites, a firewall or gateway product can be configured to use forms authentication to collect Windows credentials that are forwarded to SharePoint Server 2010. Partner employee accounts are added to the corporate directory.

- Claims authentication: This design sample incorporates the new claims authentication model. Multiple authentication providers and authentication types are implemented in a single zone. Claims authentication supports forms-based authentication, SAML token-based authentication, and Windows authentication. This design example adds partner companies using SAML token-based authentication to authenticate directly against partner directories. There are several provider options for partner employee accounts.

Use the design sample that best represents your requirements for authentication.

This article describes the design goals for the samples and explains how these goals are achieved using the logical architecture components illustrated in the samples.

# About the design samples

The design samples illustrate a corporate deployment for a fictitious company named Fabrikam, Inc. The deployment encompasses two server farms. One server farm hosts the company intranet and the partner Web site. The second farm hosts the company Web site (www.fabrikam.com). The rest of this section describes these top-level sites.

### Intranet

The corporate intranet includes the following sites:

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

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

- Emphasizes different features of SharePoint Server 2010.
- 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:

- Navigation across the applications
- Enterprise-wide search
- Shared profile data and enterprise metadata

The following diagram illustrates the three applications that make up the corporate intranet.

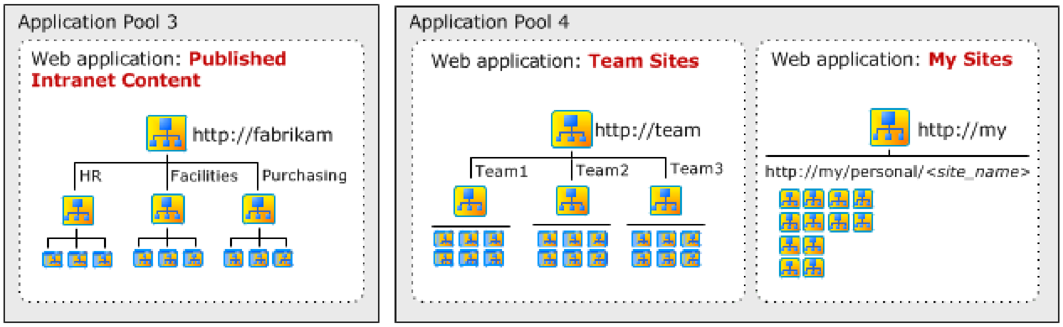
Application Pool 3
Web application: Published Intranet Content
http://fabrikam
HR
Facilities
Purchasing
Application Pool 4
Web application: Team Sites
http://team
Team1
Team2
Team3
Web application: My Sites
http://my
http://my/personal/<site_name>

The URLs within this illustration are from the classic authentication design sample.

### 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 are not allowed to access other types of content hosted on the server farm. The design for zones and service applications addresses this goal.

In the design sample, the partner Web application is hosted by the same farm that hosts the intranet content.

## Company Internet site

The company Internet site is the company's Internet presence. The content is made available to customers by configuring anonymous access with read-only permissions. Key factors that drive design choices for this application include:

- **Content isolation**: Customers cannot access any other type of content hosted on the server farm.
- **Targeted management**: Authenticated access is provided for employees who manage the Web site by performing administrative and authoring tasks.
- **Secure content authoring and publishing**: A separate site collection is hosted on Farm A in the partner Web application for authoring. This enables secure collaboration and content development with both internal and remote employees, as well as with editorial partners who specialize in Web site development or content authoring. Content publishing is configured to automatically publish the content from the authoring site collection in the first farm to the production site collection in the second farm. The following diagram illustrates the publishing process.

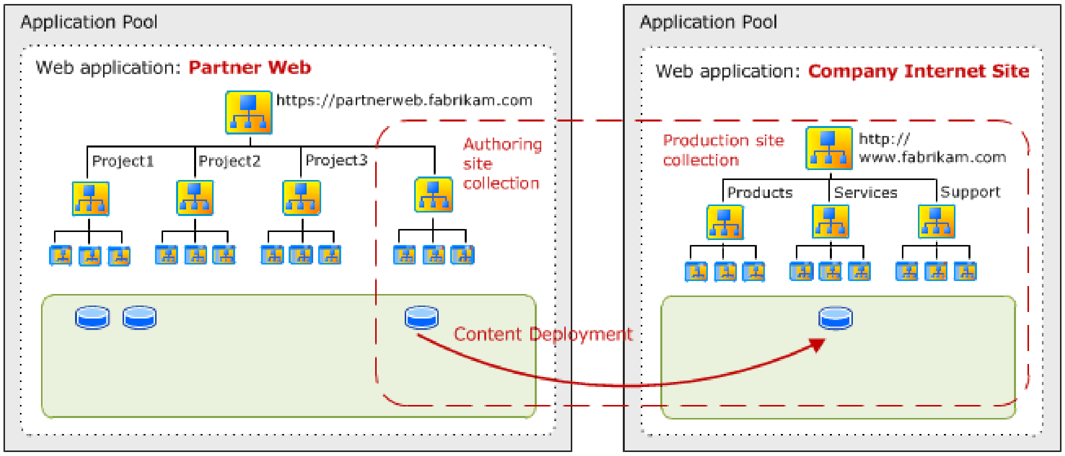
Application Pool
Web application: Partner Web
https://partnerweb.fabrikam.com
Project1
Project2
Project3
Authoring site collection
Content Deployment
Application Pool
Web application: Company Internet Site
Production site collection
http://www.fabrikam.com
Products
Services
Support

# Overall design goals

The design sample provides practical implementations of SharePoint Server 2010 features within several common types of applications. The design implementations for each of the individual applications are discussed in this article. The key design goals for the design sample include:

- Using the minimum number of server farms to host the most common types of Web sites typically required by a corporation, that is, intranet, extranet, and Internet sites.
- Creating a framework for designing an environment that can grow. Design decisions for individual applications do not prevent the addition of other applications. For example, an initial deployment might include only collaborative team sites or only the three applications that compose an intranet (team sites, My Sites, and published intranet content). By using a similar logical architecture design, you can add applications to the solution without affecting the design of the initial applications. In other words, the design does not incorporate design choices that limit the use of the environment.
- Providing 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) with different authentication providers can participate in collaboration. Also, users can only access the content they are intended to access. By following a similar logical architecture design, you create the opportunity to provide access to users in multiple locations and with different objectives. For example, your initial design might be intended only for internal employee access. However, by using a similar design you create the opportunity to also enable access to remote employees, partner employees, partner companies, and customers.
- Ensuring that the design can be used in an extranet environment. Deliberate design choices are made to ensure that the server farms 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

The design sample incorporates the use of two server farms. This section describes the licensing requirements that affect the number of server farms that are required in a corporate environment and notes the topologies of the server farms that are illustrated in the design sample.

### Licensing requirements

To host both intranet content and Internet sites, a minimum of two servers are required. This is necessary to satisfy licensing requirements.

The following two server licenses are available for SharePoint Server 2010:

- **Microsoft SharePoint Server 2010, Server License**: This is the appropriate license for collaborative intranet content. This license requires the use of Client Access Licenses (CALs). If you create sites for partner collaboration, you must ensure that you purchase the requisite number of CALs for partner employees.
- **Microsoft SharePoint Server 2010 for Internet sites**: This license is intended for Internet-facing Web sites only. This license does not require CALs. If you create sites for partner collaboration, you do not need to purchase additional CALs. However, you cannot create sites that are intended exclusively for use by your employees.

Customers who want to combine their SharePoint Server 2010 requirements under a single deployment may obtain licenses for both products, assign those licenses to the same server, and use the same running instance of the software at the same time under both licenses. However, customers must obtain CALs as required under the SharePoint Server 2010 use rights for users and devices that access content in any manner that is not permitted under the SharePoint Server 2010 for Internet Sites use rights.

For more information about farm licensing requirements, see the following resources:

- Planning for server farms (SharePoint Server 2010)
- SharePoint 2010: How to Buy (http://go.microsoft.com/fwlink/?LinkID=196728)
- SharePoint License Overview: Determining Your Needs (http://go.microsoft.com/fwlink/?LinkId=210179)
- Microsoft Volume Licensing Brief: Microsoft SharePoint Server 2010 for Internet Sites (http://go.microsoft.com/fwlink/?LinkId=210180)

Additionally, the design sample includes Microsoft Office Web Apps. Office Web Apps requires an Microsoft Office 2010 client license. In other words, if you make Office Web Apps available to partners, you must also purchase Office 2010 client licenses for them.

### Topology of the server farms

Each server farm in the design sample is composed of five servers with the following topology:

- Two front-end Web servers

- One application server
- Two database servers, either clustered or mirrored

The design sample illustrates the logical architecture of SharePoint Server 2010 by showing that:

- 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 not important to the logical architecture, except to increase capacity and performance as needed. The logical architecture can be designed independent of the server farm topology. The performance and capacity planning process will help you size the server farm to meet performance and capacity goals. For more information, see Performance and capacity management (SharePoint Server 2010).

### Scaling beyond two farms

Your business might require more than the two farms represented. 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.
- Authoring and staging sites: If published content is complex or extensive, authoring and staging might be better optimized by hosting these sites on a dedicated single-server farm using the Microsoft SharePoint Server 2010 for Internet sites license. For example, publishing content that includes tagged metadata increases the complexity of the services design sample between both the authoring farm and the published farm, including sharing the service across farms and making decisions about how the service might be shared across other types of Web applications in a multi-use farm.
- Partner sites: Security and isolation requirements might warrant a dedicated farm for partner collaboration. This creates physical isolation between internal-only content and content that is developed in collaboration with external partners.

# Users, zones, and authentication

When you create a Web application in SharePoint Server 2010, you must choose either claims-based authentication or classic-mode authentication. The authentication mode determines how accounts are used internally by SharePoint Server 2010. The following table summarizes the two approaches.

| Type of authentication | Description | Recommendations |
|---|---|---|
| Classic mode authentication | User accounts are treated by SharePoint Server 2010 as traditional Windows Active Directory accounts. The | Classic mode is recommended for upgrading environments from Microsoft SharePoint Server 2010 2007, in which forms-based |

| | | |
|---|---|---|
| | following authentication protocols are supported: Kerberos, NTLM, Basic, Digest, and anonymous.<br>Forms-based authentication is not supported.<br>Only one authentication method can be configured on a zone. | authentication is not a requirement.<br>You do not need to run user migration if you are upgrading and select classic mode authentication. |
| Claims-based authentication | User accounts are treated by SharePoint Server 2010 as claims identities. Windows accounts are automatically converted to claims identities. This mode additionally supports forms-based authentication and authentication against a trusted identity provider.<br>Multiple authentication types can be configured on a single zone. | Claims-based authentication is recommended for new SharePoint Server 2010 deployments. It is required for upgrading SharePoint Server 2010 2007 solutions that require forms-based authentication. |

The two design samples that are discussed in this article represent these two options. The following sections specifically discuss how authentication is incorporated into the two design samples.

## Classic mode authentication design sample

The design sample that uses classic mode authentication incorporates the traditional one-zone-per-type of authentication approach that was incorporated in the previous release. For this reason, this example provides a path for upgrading from SharePoint Server 2010 2007 to SharePoint Server 2010.

The one caveat is that forms-based authentication is not supported in classic mode. When using classic mode authentication, all authenticated accounts must reside within Active Directory Domain Services (AD DS). The recommendation for users who are accessing sites remotely is to use forms-based authentication on the firewall or gateway product to collect Windows credentials that are forwarded to the SharePoint farm.

The classic mode sample illustrates four different classes of users, each assigned to a different zone. Within each Web application, you can create up to five zones using one of the available zone names: Default, Intranet, Internet, Custom, or Extranet.

The following table shows the zones, users, and authentication type prescribed by the classic mode design sample.

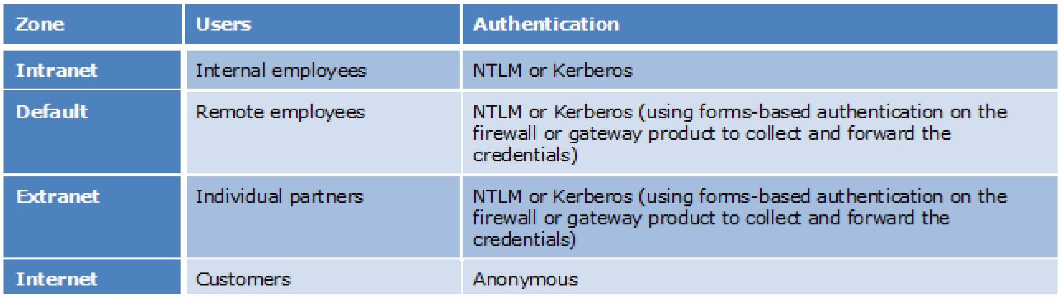

| Zone | Users | Authentication |
|---|---|---|
| Intranet | Internal employees | NTLM or Kerberos |
| Default | Remote employees | NTLM or Kerberos (using forms-based authentication on the firewall or gateway product to collect and forward the credentials) |
| Extranet | Individual partners | NTLM or Kerberos (using forms-based authentication on the firewall or gateway product to collect and forward the credentials) |
| Internet | Customers | Anonymous |

The search crawl account requires access to at least one zone using NTLM authentication. If none of the zones for users is configured to use NTLM, configure the custom zone to use NTLM authentication.

### Claims-based authentication design sample

Claims-based authentication is recommended for all new deployments of SharePoint Server 2010 and required for upgrading SharePoint Server 2010 2007 solutions that require forms-based authentication. In addition to providing the standard Windows authentication methods, claims-based authentication allows you to authenticate against other directories, such as Windows Live ID, Active Directory Federation Services 2.0, or a third-party identity provider that supports SAML tokens and the WS Federation protocol.

In the claims-based authentication design sample, claims-based authentication is used on the collaborative farm. Claims-based authentication allows multiple types of authentication to be used in the same zone. The design sample uses the Default zone for all authentication types. The following table shows the zones, users, and authentication type that are prescribed by the sample for the collaborative farm.

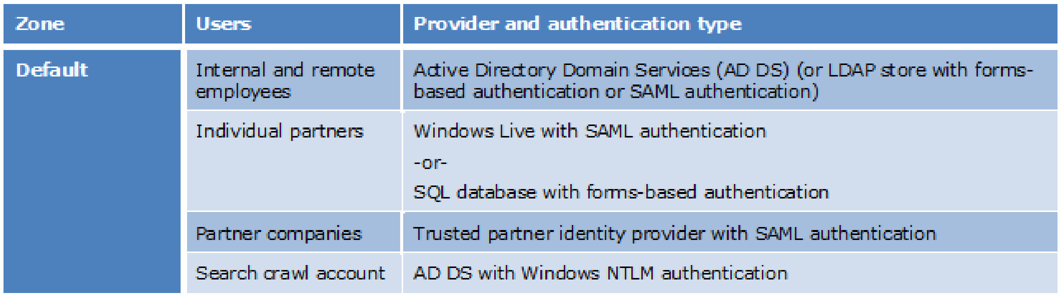

| Zone | Users | Provider and authentication type |
|---|---|---|
| Default | Internal and remote employees | Active Directory Domain Services (AD DS) (or LDAP store with forms-based authentication or SAML authentication) |
| | Individual partners | Windows Live with SAML authentication<br>-or-<br>SQL database with forms-based authentication |
| | Partner companies | Trusted partner identity provider with SAML authentication |
| | 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) that can be used both internally and externally. Active Directory accounts are used. If needed, LDAP can be used with either forms-based authentication or SAML, which requires additional configuration.

In the design sample, the partner Web application represents an extranet site that is accessible by partner employees and partner companies. Using claims-based authentication in this scenario requires that trust be configured with one or more external secure token service (STS). This can be provided using either one of the following approaches:

- The SharePoint farm can be configured to trust an external STS, such as the STS associated with Windows Live (for authenticating individual partners) or the STS that resides in a partner company (for authenticating directly against the partner directory).
- The STS inside the corporate environment can be configured to trust an external STS. This relationship must be established explicitly by administrators in the two organizations. In this scenario, the SharePoint farm is configured to trust only the STS that resides in its own corporate environment. This internal STS verifies the token it receives from the external STS, and then issues a token that allows the partner user to access the SharePoint farm. This is the recommended approach.

An alternative to implementing a claims-based environment to authenticate partners is to use forms-based authentication and manage these accounts using a separate store, such as a database.

For more information about implementing a claims-based authentication environment, see the following white paper: Claims-based Identity for Windows: An Introduction to Active Directory Federation Services 2.0, Windows CardSpace 2.0, and Windows Identity Foundation (http://go.microsoft.com/fwlink/?LinkId=196776).

In the claims-based authentication design sample, the published farm is set up to use classic mode authentication. An alternative is to use claims-based authentication for the published farm as well and implement a separate zone for anonymous users. The important element of the design is to use a separate zone for anonymous users to create isolation between the read-only environment and the read-write environment, regardless of which authentication mode is implemented. The following table shows the zones, users, and authentication type that are illustrated for the published farm.

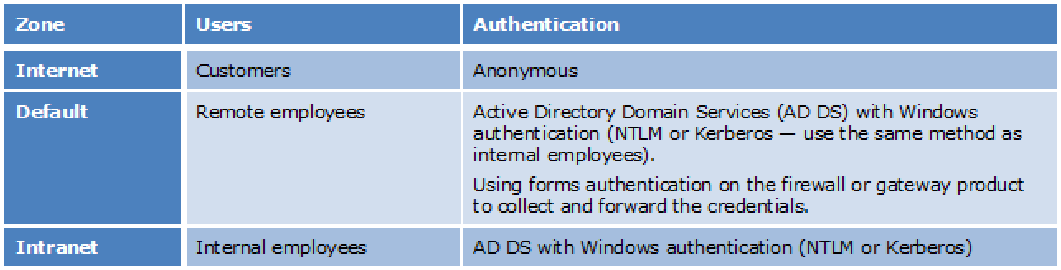

| Zone | Users | Authentication |
|---|---|---|
| **Internet** | Customers | Anonymous |
| **Default** | Remote employees | Active Directory Domain Services (AD DS) with Windows authentication (NTLM or Kerberos — use the same method as internal employees).<br>Using forms authentication on the firewall or gateway product to collect and forward the credentials. |
| **Intranet** | Internal employees | AD DS with Windows authentication (NTLM or Kerberos) |

Again, the search crawl account requires access to at least one zone using NTLM authentication. NTLM authentication can be added to a claims-authentication zone, if needed. In classic-mode, if none of the zones for users is configured to use NTLM, configure the custom zone to use NTLM authentication.

## 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 Server 2010 places the following requirements on how the Default zone is configured:

- 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 is sent with links from the Default zone. These include e-mail to owners of sites that are approaching quota limits. Consequently, users who receive these type of e-mails and alerts must be able to access links through the Default zone. This is especially important for site owners.
- Host-named site collections are only available through the Default zone. All users who are intended to access host-named site collections must have access through the Default zone.

### Configuring zones for an extranet environment

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

- User requests can be initiated from several different networks. In the design sample, 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 different Web applications. Additionally, internal and remote employees can potentially contribute to and administer content across all of the Web applications: intranet, Partner Web, and the company Internet site.

In an extranet environment, ensure that the following design principles are followed:

- Configure zones across multiple Web applications to mirror each other. The configuration of authentication and the intended users should be the same. However, the policies associated with zones can differ across Web applications. For example, ensure that the intranet zone is used 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.
- Configure alternate access mappings appropriately and accurately for each zone and each resource. Alternate access mappings are automatically created when you create the zone. However, SharePoint Server 2010 can be configured to crawl content in external resources, such as a file share. Links to these external resources must be created manually for each zone by using alternate access mappings.

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

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

# Services

The services architecture illustrated shows the most complex option for deploying services across the three different types of sites: Intranet, partner Web, and the company Internet site. Dedicated and partitioned services are deployed for the partner Web site. A separate instance of the Managed Metadata service application is deployed for exclusive us by the authoring site collection and published Internet site.

A much simpler alternative is to deploy one set of service applications and to share each service, as needed, across the sites. This architecture relies on security trimming to show only content that users have access to. The following diagram illustrates this simpler approach.

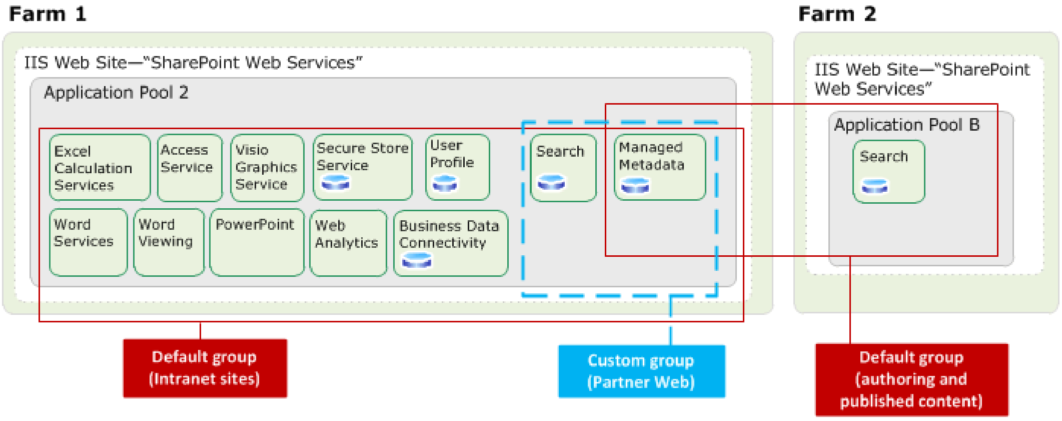
Farm 1
IIS Web Site—"SharePoint Web Services"
Application Pool 2
Excel Calculation Services
Access Service
Visio Graphics Service
Secure Store Service
User Profile
Search
Managed Metadata
Word Services
Word Viewing
PowerPoint
Web Analytics
Business Data Connectivity
Farm 2
IIS Web Site—"SharePoint Web Services"
Application Pool B
Search
Default group (Intranet sites)
Custom group (Partner Web)
Default group (authoring and published content)

The primary design decision for deploying service applications is how broadly to spread the organization taxonomy. Services architecture can be simplified by sharing managed metadata, user profile, and search across all Web apps and relying on security trimming to manage access to content. In the simplified architecture described in this article, one instance of the Managed Metadata service is shared across all sites. However, with this configuration all users have access to the corporate taxonomy. 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.

### Partner Web site

For the partner Web site (custom group on Farm 1), the minimum services prescribed by the design sample are Search and Managed Metadata. If you add Office Web Apps to the group of services used by the partner Web site, ensure that you have the appropriate licenses for all users of this site, including partners. The User Profile service application is not included by the design sample to prevent partner users from browsing people data in the organization.

In the simplified architecture, partners have access to the entire corporate taxonomy and can browse people data in the organization. However, search limits results to sites and content that partners have access to.

If your partner sites require content isolation between projects, deploying dedicated and partitioned service applications is a good choice, as illustrated in the design sample. This increases the complexity of the services architecture but ensures that partners do not have access to metadata associated with the Intranet content or even other projects within the partner Web site.

### Company Internet site

In the simplified design architecture, the corporate Managed Metadata service application is also shared with the published Internet site. In the design sample, a dedicated instance of the Managed Metadata service application is deployed on the collaboration farm for exclusive use by the authoring site collection and the published farm.

If the published farm is anonymous and read-only, there is no risk of exposing managed metadata that is not associated with the published content. Anonymous users have access only to the content that is published and cannot submit ratings or create other types of metadata.

Sharing the Managed Metadata service application across the organization (as illustrated in the simplified architecture in this article) allows authors to utilize the corporate taxonomy. In contrast, deploying a dedicated instance of the service for authoring and publishing (illustrated in the design sample) ensures that managed metadata is isolated.

A dedicated instance of the Search service application is deployed to the farm hosting the company Internet site. This is the recommended configuration for a published Internet-facing site.

# Authoring and publishing alternatives

For the company Internet site, the design sample illustrates a publishing process that includes using the content deployment feature to move content from an authoring site collection to the publishing farm. A simpler alternative to this approach is to author directly on the publishing farm. This is commonly referred to as authoring in production.

Authoring on the production environment greatly simplifies the solution by consolidating services on one farm and removing the need for content deployment. The design sample includes the additional zones that can be used by authors to work securely without impacting anonymous users. Be sure to block incoming anonymous access on the port associated with the zones used by authors. If your site has less than 500 writes per hour of authoring activity, it is unlikely that performance of the published site will be impacted when authoring on the production environment.

SharePoint Server 2010 includes publishing features that can be used in this scenario to ensure that content is not exposed to anonymous users until it is ready. For more information, see the following articles:

- Schedule the start and end date for a published page (http://go.microsoft.com/fwlink/?LinkId=196777)
- Approve or reject a pending submission (http://go.microsoft.com/fwlink/?LinkId=196778)
- Set permissions for publishing (http://go.microsoft.com/fwlink/?LinkId=196779)

# Administration sites

In the design sample, the Central Administration site for each server farm is hosted on an application server. 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 Central Administration site remains available.

The load-balanced URLs for administration sites are not mentioned in the design sample or in this article. Recommendations include the following:

- If port numbers are used in administrative URLs, use non-standard ports. Port numbers are included in URLs by default. While port numbers are typically not used in customer-facing URLs, 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 Central Administration site 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 mitigates an exploit on one site that provides an opportunity for an attacker to inject code onto the server to attack other sites.

Practically speaking, consider using a dedicated application pool for each of the following scenarios:

- To separate authenticated content from anonymous content.
- To isolate applications that store passwords for and interact with external business applications (although the Secure Store Service can be used for this purpose instead).
- To isolate applications where users have great liberty to create and administer sites and to collaborate on content.

The design sample uses application pools in the following way:

- Each administration site is hosted in a dedicated application pool. This is a requirement of SharePoint 2010 Products.
- Intranet content is divided into two different application pools. Collaborative content (My Sites and team sites) is hosted in one application pool. The published intranet content is hosted in a separate application pool. This configuration provides process isolation for the published intranet content in which business data connections are more likely to be used.
- The partner Web application is hosted in a dedicated application pool.
- The company Internet site is hosted in a dedicated application pool on the second farm. If this farm were also to host content for partner collaboration, these two types of content (Internet and partner) would be hosted in two different application pools.

# Web applications

A Web application is an IIS Web site that is created and used by SharePoint 2010 Products. Each Web application is represented by a different Web site in IIS.

Generally speaking, use dedicated Web applications to:

- **Separate anonymous content from authenticated content.** In the design sample, the company Internet site is hosted in a dedicated Web application and application pool.
- **Isolate users.** In the design sample, the partner Web site is hosted in a dedicated Web application and application pool to ensure that partners do not have access to the intranet content.
- **Enforce permissions.** A dedicated Web application provides the opportunity to enforce permissions by policies by using the Policy for Web Application page in Central Administration. For example, you can create a policy on the company Internet site to explicitly deny write access to one or more groups of users. Policies for a Web application are enforced regardless of permissions configured on individual sites or documents within the Web application.
- **Optimize performance.** Applications achieve better performance if they are placed 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 allows you to restore more critical content before restoring My Site content. In the design sample, My Sites are placed in a separate Web application to enable administrators to more aggressively manage growth compared to other applications.

# Site collections

Site collections bridge logical architecture and information architecture. The design goals for site collections in the design sample are to satisfy requirements for URL design and to create logical divisions of content.

To satisfy the requirements for URL design, each Web application includes a single root-level site collection. Managed paths are used to 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.

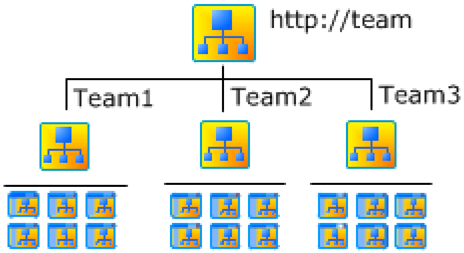

Given the requirement for a root-level site collection, the design decisions revolve around the second tier of site collections. The design sample incorporates choices based on the nature of the application.

### 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, each division's content is hosted in a separate site collection. This provides the following advantages:

- Each division can manage and permission their content independently.
- Each division's content can be stored in a dedicated database.

The disadvantages of using multiple site collections include:

- Master pages, page layouts, templates, Web Parts, and navigation cannot be shared across site collections.
- More effort is required to coordinate customizations and navigation across site collections.

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

### My Sites

My Sites have distinct characteristics and the recommendations for deploying My Sites are straightforward. In the design sample, the My Site application 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 inherit the My Site Host template. The user name is appended to the URL in the form http://my personal/*username*. The following illustration illustrates My Sites.

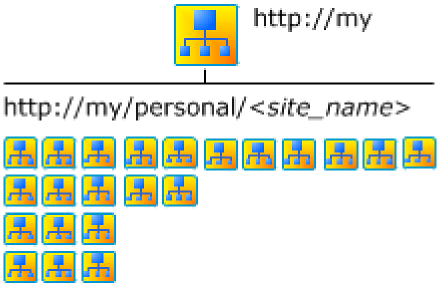

**Team sites**

You can use either of the following two approaches for designing site collections within 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. However, there are many disadvantages to this approach, including:
    - You lose the opportunity to implement a thoughtful taxonomy.
    - The application can become difficult to manage.
    - Sites are easily abandoned.
    - 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, site collections are created by a SharePoint administrator. After the site collection is created, teams can create sites within the site collection based on their needs. 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.

The design sample incorporates the second approach, which results in a similar site collection hierarchy for team sites as for 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 shows suggestions for different types of organizations.

| **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.<br>• 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.<br>• Create a site collection for each category of research projects. |

| Type of organization | Suggested site collection taxonomies |
|---|---|
| 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.<br>• 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 within 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 have access to only the project they work on.
- Permissions are managed by site owners.
- 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. In this way, the following benefits occur:

- Individual site collections provide the appropriate level of isolation between projects.
- Self-service site creation can be implemented.

Because the partner Web application also hosts the site collections for developing content for the company Internet site, separate site collections are created for authoring and staging.

### Company Internet site

The company Internet site incorporates a single root-level site collection. All sites beneath this site collection are subsites. This structure simplifies URLs for pages within the site. The following diagram illustrates the architecture of the company Internet site.

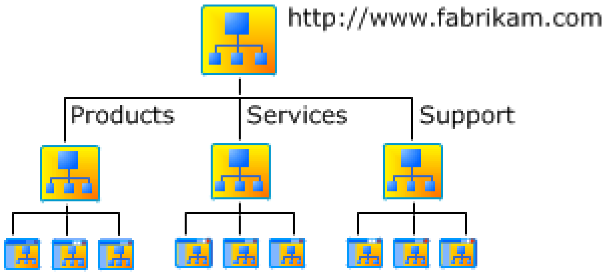

## Content databases

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 size warning thresholds are reached. 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 can be managed 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 performing the following steps:
  a. Within the Web application, create the database and set the database status to **Ready**.
  b. 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.
  c. Create the site collections. They are automatically added to the database.
  d. Set the status of all other databases back to **Ready**.

### Published intranet content

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

### My Sites

For My Sites, the design sample achieves 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 how much additional space 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 into the database.

The design sample provides the following example size settings based on a target database size of 200 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 alternately to the new database and the existing database until the maximum number of sites for one of the databases is reached.

### Team sites

For team sites, again the design sample achieves 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 with teams that 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. In the design sample, however, Partner Web also hosts the authoring site collection for the company Internet site. Consequently, database design incorporates both approaches:

- Authoring site collection is hosted in a dedicated database.
- Database and size settings are configured to manage all other sites and databases.

Because Partner Web is hosted in a dedicated Web application, you can create size limits that are more appropriate to the types of sizes that are created. The design sample provides 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

### Company Internet site

By using a single site collection in the design of the company Internet site, you use a single database for this Web application.

# Zones and URLs

The design sample illustrates how to coordinate URLs across multiple applications within a corporate deployment.

### Design goals

The following goals influence design decisions for URLs:

- URL conventions do not limit the zones through which content can be accessed.
- 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.

### Design principles

To achieve these design goals, the following design principles are applied:

- Host-named site collections are not used. Note that host-named site collections are different from IIS host headers. Host-named site collections do not work with the alternate access mappings feature. The alternate access mappings feature is required to access the same content through multiple domain URLs. Consequently, when host-named sites are used, these sites are available only through the Default zone.
- Each application incorporates a single root site collection. This is a requirement for using alternate access mappings. If multiple root site collections are required within a Web application and you expect to use only the Default zone for user access, host-named site collections are a good option.
- For the applications that include multiple high-level site collections, in which each site collection represents a top-level team or project (for example, team sites), the design sample incorporates managed paths. Managed paths provide greater control over URLs for these types of sites.

### Design tradeoffs

Meeting the design goals results in some tradeoffs, including the following:

- URLs are longer.
- Host-named site collections are not used.

## Designing load-balanced URLs

When you create a Web application, you must choose a load-balanced URL to assign to the application. Additionally, you must create a load-balanced URL for each 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 application and each zone within each application requires a unique URL across the design sample.

### Intranet

Each of the three Web applications that compose the intranet requires a unique URL. In the design sample, the target audience for the intranet content is internal employees and remote employees. In the claims authentication design sample, employees use the same URLs for each of these applications regardless of whether they are on site or remote. While this approach adds a layer of security to the SharePoint design (all traffic is SSL), this approach requires either routing internal traffic through the firewall or gateway product along with remote traffic or setting up a split DNS environment to resolve internal requests within the internal network.

For the classic authentication design sample, the URLs are different for internal and remote employees. The following table shows the URLs for internal and remote employees to access each application in the classic authentication design sample.

| Application | Internal employee URL | Remote employee URL |
|---|---|---|
| Published intranet content | http://fabrikam | https://intranet.fabrikam.com |
| Team sites | http://teams | https://teams.fabrikam.com |
| My Sites | http://my | https://my.fabrikam.com |

### Partner Web site

In the design sample, the partner Web site is accessed by internal employees, remote employees, and partner employees. In the claims authentication design sample, each of these users enters the same URL, regardless of the authentication method. In the classic authentication design sample, each different type of user enters a different URL. Although both remote employees and partner employees access the partner Web site externally using SSL (HTTPS), each group requires a different URL to apply the benefits of using separate zones—that is, different authentication methods and different zone policies. The following table shows the URLs that internal employees, remote employees, and partners use to access the partner Web site, as shown in the classic authentication design sample.

| Zone | URL |
|---|---|
| Internal employee URL | http://partnerweb |
| Remote employee URL | https://remotepartnerweb.fabrikam.com |
| Partner URL | https://partnerweb.fabrikam.com |

**Company Internet site**

The company Internet site is a public site and can be accessed by any user by using the default URL, http://www.fabrikam.com. The policies of the Internet zone are applied (that is, anonymous access and deny write).

However, to support administration and authoring tasks on the public site, the design sample incorporates URLs for internal and remote employees. You can use policies for these zones ensure appropriate access to targeted security groups for authoring and maintenance tasks. Both the classic authentication and the claims authentication design samples use the same approach for this farm. The following table shows the URLs for each zone.

| Zone | URL |
|---|---|
| Internal employee URL | http://fabrikamsite |
| Remote employee URL | https://fabrikamsite.fabrikam.com |
| Customer URL | http://www.fabrikam.com |

## Using explicit and wildcard inclusions for URL paths

By defining managed paths, you can specify which paths in the URL namespace of a Web application 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 created 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://fabrikam/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 created by using this method is http://my/personal/user1.

The design sample incorporates the use of both types as described in the following sections.

### Explicit inclusions: Published intranet content

In the design sample, the published intranet Web application incorporates the use of explicit inclusions.

#### Published intranet content

Within the published intranet content Web application, an explicit inclusion is used for each subsite, for example, HR, Facilities, and Purchasing. Each of these site collections can be associated with a different content database, if needed. 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 recommended limit for sites created using an explicit inclusion is about 20. If your organization requires a greater number of site collections, then consider using a wildcard inclusion or host-named site collections.

In the classic authentication design sample, the use of explicit inclusions results in the URLs shown in the following table.

| Internal employee (Intranet zone) | Remote employee (Default zone) |
|---|---|
| http://fabrikam | https://intranet.fabrikam.com |
| http://fabrikam/hr | https://intranet.fabrikam.com/hr |
| http://fabrikam/facilities | https://intranet.fabrikam.com/facilities |
| http://fabrikam/purchasing | https://intranet.fabrikam.com/purchasing |

In this example, the root site collection, http://fabrikam, 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 a lot of 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.

In the classic authentication design sample, the use of wildcard inclusions results in the URLs shown in the following table.

| Internal employee (Intranet zone) | Remote employee (Default zone) |
|---|---|
| http://teams/sites/Team1 | https://teams.fabrikam.com/sites/Team1 |
| http://teams/sites/Team2 | https://teams.fabrikam.com/sites/Team2 |
| http://teams/sites/Team3 | https://teams.fabrikam.com/sites/Team3 |

In this example, the root site collection, http://team, does not necessarily host content for users.

### My Sites

My Sites offer self-service site creation. When a user browsing 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.

In the classic authentication design sample, this results in URLs of the format shown in the following table.

| Internal (Intranet zone) | Remote employee (Default zone) |
|---|---|
| http://my/personal/user1 | https://my.fabrikam.com/personal/user1 |
| http://my/personal/user2 | https://my.fabrikam.com/personal/user2 |
| http://my/personal/user3 | https://my.fabrikam.com/personal/user3 |

### Partner Web application

T partner Web application is designed to allow employees to easily create secure sites for collaboration with external partners. To support this goal, self-service site creation is allowed.

In the classic authentication design sample, the partner Web application includes a wildcard inclusion named /sites (http://partnerweb/sites). This results in URLs of the format shown in the following table.

| Internal employee (Intranet zone) | Remote employee (Default zone) |
|---|---|
| http://partnerweb/sites/project1 | https://remotepartnerweb.fabrikam.com/sites/project1 |
| http://partnerweb/sites/project2 | https://remotepartnerweb.fabrikam.com/sites/project2 |
| http://partnerweb/sites/project3 | https://remotepartnerweb.fabrikam.com/sites/project3 |

Partner contributors can access partner Web sites using the URLs listed in the following table.

| Partner (Extranet zone) |
|---|
| https://partnerweb.fabrikam.com/sites/project1 |
| https://partnerweb.fabrikam.com/sites/project2 |
| https://partnerweb/fabrikam.com/sites/project3 |

The exception for the partner Web application, as illustrated in the design samples, is the collection dedicated to authoring the content for the company Internet site. For this site collection, an explicit inclusion is used. This provides an example of using both explicit inclusions and wildcard inclusions in the same Web application.

# Zone policies

You can create a policy for a Web application to enforce permissions at the Web application level. A policy can be defined for the Web application in general or just for a specific zone. A policy enforces permissions on all content within the Web application or the 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 re-crawl sites to pick up the new permissions.

Policies are not used in the claims authentication design sample for the collaborative farm where multiple types of authentication are enabled on a single zone. Policies are implemented in the classic authentication design sample and on the published farm of the claims authentication design sample where windows authentication is prescribed. On the published farm, the use of policies adds an additional layer of security between anonymous users and users who have access to manage sites.

The design samples provide examples of several policies to accomplish the following:

- Deny write access to published content.
- Ensure authors and testers have appropriate access to published content.

# Plan for host-named site collections (SharePoint Server 2010)

In this article:

- About host-named site collections
- About host headers
- Create a host-named site collection
- Programmatically create a host-named site collection
- Use managed paths with host-named site collections
- Expose host-named sites over HTTP or SSL
- Configure SSL for host-named site collections
- Use host-named site collections with off-box SSL termination

Microsoft SharePoint Server 2010 supports both path-based and host-named site collections. The primary difference between path-based and host-named site collections is that all path-based site collections in a Web application share the same host name (DNS name), and each host-named site collection in a Web application is assigned a unique DNS name.

Path-based site collections provide a corporate hosting solution with all site collections sharing the same host name of the Web application. In a path-based deployment, you can have a single site collection at the root of the Web application and additional site collections under managed paths within the Web application.

Host-named site collections provide a scalable Web hosting solution with each site collection assigned to a unique DNS name. In a Web hosting deployment, each host-named site collection has its own vanity host name URL, such as http://customer1.contoso.com,  http://customer2.contoso.com, or http://www.customer3.com.

SharePoint Server 2010 provides two significant improvements to host-named site collections: the ability to use managed paths with host-named site collections, and the ability to use off-box SSL termination with host-named site collections.

## About host-named site collections

Web hosters provide customers with Web server space to host their own Web sites. In a path-based SharePoint Server 2010 environment, these sites would typically be assigned to http://www.contoso.com/sites/customer1,  http://www.contoso.com/sites/customer2, and so on. However, Web hosting customers typically want to have their Web sites available at a vanity domain name, such as http://customer1.contoso.com, http://customer2.contoso.com, and so on.

One way to support this customer request is to provide each customer with their own Web application and assign the customer's unique DNS name to the Web application. However, SharePoint Server

2010 Web applications do not scale as well as SharePoint Server 2010 site collections. SharePoint Server 2010 supports host-named site collections as an alternative to creating individual Web applications for each customer. Host-named site collections can scale to thousands of site collections because they can all exist within a single Web application and still offer vanity naming capability.

Because host-named site collections have a single URL, they do not support alternate access mappings and are always considered to be in the Default zone. If you need to support site collections responding to multiple host-name URLs, consider using path-based site collections with alternate access mappings instead of host-named site collections.There are several additional configuration options to consider when provisioning a new SharePoint Server 2010 site. Specifying the appropriate site template during site creation will determine which preconfigured Web parts and other user interface elements are available on the new site. In a hosting scenario, you will probably want to select either a team site template (value of "STS#0" when creating the site) or a blank site with no Web parts or prebuilt lists (value of "STS#1"). Also consider specifying site quotas on each newly provisioned site collection.

# About host headers

Host headers refer to the portion of the HTTP protocol that tells the Web server the DNS name of the site that the client is connecting to. You can apply host headers at two different levels in SharePoint Server 2010:

- The Web application (IIS Web site) level
- The site collection level

It's important to understand the distinction between these two levels. Host headers at the IIS Web site level are only intended for path-based site collections. Host headers at the site collection level are only intended for host-named site collections. In most cases, applying a host header binding at the IIS Web site level makes it impossible to access host-named site collections through the IIS Web site. This is because IIS will not respond to requests for host names that differ from the host header binding.

Path-based site collections and host-named site collections can co-exist in the same Web application and can exist in multiple Web applications. To ensure that both types of site collections are accessible to users, do not put host header bindings on the IIS Web site assigned to the Default zone of your Web application, if you have host-named site collections in that Web application. You can apply host header bindings to the IIS Web sites in the other zones of your Web application. This enables you to use the Default zone with host-named site collections while allowing you to use alternate access mapping functionality in the other zones for path-based site collections.

You can manually modify host header bindings on the IIS Web site from the IIS Manager, but this is not recommended. Any changes you make using the IIS Manager will not be recorded in SharePoint Server 2010. If SharePoint Server 2010 tries to provision an IIS Web site on another computer in the farm for the same Web application and zone, the original host header binding is used instead of the modified binding. If you want to modify an existing binding for an IIS Web site, remove the Web application from the zone and then re-extend the Web application into the zone with the binding you want to use.

# Create a host-named site collection

You must use Windows PowerShell to create a host-named site collection. You cannot use the SharePoint Server 2010 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:

1. To create a host-named site collection using Windows PowerShell, verify that you meet the following minimum requirements: See Add-SPShellAdmin.
2. On the **Start** menu, click **All Programs**.
3. Click **Microsoft SharePoint 2010 Products**.
4. Click **SharePoint 2010 Management Shell**.
5. From the Windows PowerShell command prompt (that is, PS C:\>), type the following:

```
New-SPSite http://host.header.site.url -OwnerAlias DOMAIN\username -

HostHeaderWebApplication http://servername
```

This creates a host-named site collection with the URL `http://host.header.site.url` in the SharePoint Server 2010 Web application with the URL `http://servername`.

# Programmatically create a host-named site collection

In addition to using the Windows PowerShell to create host-named sites, you can use the SharePoint Server 2010 object model. The following code sample creates the host-named site collection with the URL `http://host.header.site.url` in the SharePoint Server 2010 Web application with the URL `http://servername`:

```
SPWebApplication webApp = SPWebApplication.Lookup(new

Uri("http://www.contoso.com"));

SPSiteCollection sites = webApp.Sites;

SPSite Site = null;

Site = sites.Add("http://hoster.contoso.com", "Site_Title",

"Site_Description", 1033, "STS#0", "contoso\owner",

"Owner_Display_Name", "Owner_Email", "contoso\secondaryowner,

"Secondary_Owner_Display_Name", "Secondary_Owner_Email", true);
```

SharePoint Server 2010 ships with a set of Web services for various user and administrative tasks. One of these administrative tasks is creating a new site collection. The **CreateSite** Web service method does not support the creation of host-named site collections. A workaround for this issue is to write a Web service that wraps the API sample code.

# Use managed paths with host-named site collections

SharePoint Server 2010 adds support for 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 different from managed paths for path-based site collections. Managed paths for host-named site collections do not apply to path-based site collections; nor do managed paths for path-based site collections apply to host-named site collections. Managed paths created for host-named site collections are available to all host-named site collections within the farm regardless of which Web application the host-named site collection is in. You must create a root host-named site collection for a host name before you can create a managed path host-named site collection for that host name.

You can create a managed path for use with host-named site collections by using the Windows PowerShell `New-SPManagedPath` cmdlet with the `-HostHeader` parameter, as shown in the following example:

```
New-SPManagedPath pathname -HostHeader
```

A host-named site collection created at a managed path is shown in the following example:

```
New-SPSite http://host.header.site.url/pathname/sitename -OwnerAlias DOMAIN\username -
HostHeaderWebApplication http://servername
```

# Expose host-named sites over HTTP or SSL

Host-named site collections will use the same protocol scheme as the public URL in the Default zone of their Web application. If you wish to provide the host-named site collections in your Web application over HTTP, ensure that the public URL in the Default zone of your Web application is an HTTP-based URL. If you wish to provide host-named site collections in your Web application over SSL, ensure that the public URL in the Default zone of your Web application is an HTTPS-based URL.

Unlike an earlier version, SharePoint Server 2010 does not support a host-named site collection using both HTTP- and SSL-based URLs simultaneously. If some host-named site collections need to be available over HTTP while other host-named site collections need to be available over SSL, separate the host-named site collections into two different Web applications dedicated for that type of access. In this scenario, HTTP host-named site collections should be in a Web application dedicated for HTTP access and SSL host-named site collections should be in a Web application dedicated for SSL access.

# Configure SSL for host-named site collections

In hosting scenarios, hosters can configure a single Web application with SSL and then create multiple host-named site collections within that Web application. To browse to a site over SSL, a server certificate has to be installed and assigned to the IIS Web site. Each host-named site collection in a Web application will share the single server certificate assigned to the IIS Web site.

Hosters need to acquire a wildcard certificate or subject alternate name certificate and then use a host-named site collection URL policy that matches that certificate. For example, if a hoster acquires a *.contoso.com wildcard certificate, the hoster has to 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 customers require unique second-level domain names for their sites, the hoster has to create multiple Web applications rather than multiple host-named site collections.

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

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

Because SharePoint Server 2010 uses the public URL in the Default zone of the Web application to determine whether host-named site collections will be rendered as HTTP or SSL, you can now use host-named site collections with off-box SSL termination. There are 3 requirements to use SSL termination with host-named site collections:

- The public URL in the Default zone of the Web application must be an HTTPS-based URL.
- 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.

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 Server 2010 will render links of its host-named site collections in that Web application using HTTPS instead HTTP.

# Hosted environments (SharePoint Server 2010)

In this section:

- Model: Hosting architectures for SharePoint Server 2010
- White paper: SharePoint 2010 for hosters (SharePoint Server 2010)

# Model: Hosting architectures for SharePoint Server 2010

This model summarizes the support for hosting environments and illustrates common hosting architectures. Before learning about hosting environments, it is important to understand the services architecture. For more information, see Services architecture planning (SharePoint Server 2010).

**Hosting Environments in SharePoint 2010 Products**

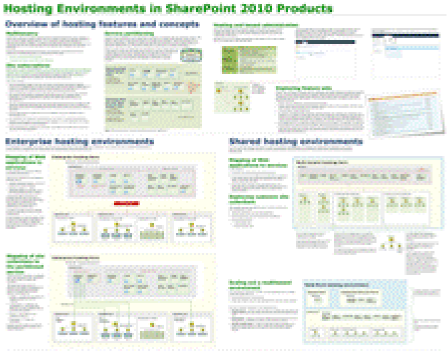
Hosting Environments in SharePoint 2010 Products
Overview of hosting features and concepts
Enterprise hosting environments
Shared hosting environments

Visio (http://go.microsoft.com/fwlink/?LinkID=167084)

PDF (http://go.microsoft.com/fwlink/?LinkID=167086)

XPS (http://go.microsoft.com/fwlink/?LinkID=167085)

# White paper: SharePoint 2010 for hosters (SharePoint Server 2010)

This white paper provides the hosting community with an overview of Microsoft SharePoint Server 2010 and the underlying Microsoft SharePoint Foundation components, and provides detailed architectural guidance to support the multi-tenancy requirements of hosters. Although this document can be used by anyone to gain an understanding of many key features in SharePoint 2010 Products, it focuses on how these new features affect and support the hosting community.

SharePoint 2010 for hosters (http://go.microsoft.com/fwlink/?LinkId=190783)

# Virtualization planning (SharePoint Server 2010)

This section contains articles that are designed to help you plan and implement a server virtualization solution for Microsoft SharePoint Server 2010 server farms.

In this section:

- Virtualization support and licensing (SharePoint Server 2010)
- Hyper-V virtualization requirements (SharePoint Server 2010)
- Plan virtual architectures (SharePoint Server 2010)
- Plan for virtualization (SharePoint Server 2010)
- Capacity management and high availability in a virtual environment (SharePoint Server 2010)

# Virtualization support and licensing (SharePoint Server 2010)

This article provides support and licensing information for using server virtualization technologies to deploy SharePoint 2010 Products in a virtual environment.

## SharePoint 2010 Products support for virtualization

All elements of Microsoft SharePoint Server 2010 are fully supported when deployed in a Windows Server 2008 Hyper-V technology environment. In addition, any related or required supporting technologies are also supported.

**Note:**
Support for SharePoint Server 2010 virtualization includes third-party virtualization technologies that are hosted or hardware-based, and certified by Microsoft. For more information about certification and participating vendors, see the Server Virtualization Validation Program (SVVP) (http://go.microsoft.com/fwlink/?LinkId=125649).

## Server virtualization using Hyper-V technology

Beginning withWindows Server 2008, server virtualization using Hyper-V has been an integral part of the operating system. Hyper-V is available with all editions of the operating system, as well as with Microsoft Hyper-V Server 2008.

We recommend using Windows Server 2008 R2 or Microsoft Hyper-V Server 2008 R2 as virtualization servers for your SharePoint 2010 Products deployment. These releases provide:

- Added capabilities, such as increased virtual processor support and increased memory support for virtual machines.
- Performance improvements, such as improved virtual hard drive performance and network adapter performance.

For more information, see What's New in Hyper-V in Windows Server 2008 R2 (http://go.microsoft.com/fwlink/?LinkID=155234).

## Operating system environment (OSE) licensing

Before you start planning for virtualization you need to determine the licensing requirements for your virtualization environment. Two types of operating system environments (OSEs) exist:

- One physical operating system environment
- One or more virtual operating system environment(s)

A virtual operating system environment is configured to run on a virtual (or otherwise emulated) hardware system. Use of technologies that create virtual OSEs does not change the licensing requirements for the operating system and any applications running in the OSE.

The Windows Server operation system licensing model for physical multicore processor systems is based on the number of physical processors installed on the hardware. This model extends to virtual processors configured for a virtual machine running on a virtualization server. For licensing purposes, a virtual processor is considered to have the same number of threads and cores as each physical processor on the underlying physical hardware system.

For more information about licensing requirements:

- Licensing Microsoft Server Products in Virtual Environments (http://go.microsoft.com/fwlink/?LinkId=187741)

  This white paper gives an overview of Microsoft licensing models for the server operating system and server applications under virtual environments.

- Windows Server Virtualization Calculators (http://go.microsoft.com/fwlink/?LinkId=187742)

  The Windows Server Virtualization Calculators provide two ways to estimate the number and cost of Windows Server Standard Edition, Enterprise Edition, and Datacenter Edition licenses needed for your virtualization scenarios to help you determine the most cost-effective edition of Windows Server.

**Note:**
Although Microsoft Hyper-V Server 2008 R2 does not require a license for the virtualization server, licensing requirements must be met for the virtual OSEs.

# SharePoint 2010 Products licensing

Every element of a SharePoint farm that is installed on a virtual machine must comply with the licensing requirements for SharePoint Server 2010 as well as related and supporting technologies.

# Hyper-V virtualization requirements (SharePoint Server 2010)

This article provides hardware and software requirements for using hardware-based virtualization. Although Windows Server 2008 Hyper-V technology is the focal point of this document, the basic hardware requirements for enabling hardware-based virtualization also apply to third-party virtualization technologies that are certified by Microsoft.

## Hardware

The requirements for hardware-based virtualization are as follows:

- Hardware-assisted virtualization, which is available in processors that include a virtualization option—specifically processors with Intel Virtualization Technology (Intel VT) or AMD Virtualization (AMD-V) technology.
- Hardware-enforced Data Execution Prevention (DEP) is available and enabled.

You can use one of the following tools to determine if the processor on an existing server supports Hyper-V:

- AMD Hyper-V Compatibility Check Utility (.zip file) (http://go.microsoft.com/fwlink/?LinkId=150561)
- Intel Processor Identification Utility (Windows Version) (http://go.microsoft.com/fwlink/?LinkId=150562)

## Software

One of the following Microsoft products is required for Hyper-V:

- Windows Server 2008 (all editions of Windows Server 2008, except for Windows Server 2008 for Itanium-Based Systems, Windows Web Server 2008, and Windows Server 2008 Foundation)
- Microsoft Hyper-V Server 2008
- Windows Server 2008 R2 (all editions of Windows Server 2008 R2, except for Windows Server 2008 R2 for Itanium-Based Systems, Windows Web Server 2008 R2, and Windows Server 2008 R2 Foundation)
- Hyper-V Server R2

We recommend Windows Server 2008 R2 for virtualization servers because of the many improvements introduced for Hyper-V, such as:

- Live migration to move a running virtual machine from one cluster node to another
- Significant gains in performance and scalability
- Enhanced processor support
- Enhanced virtual machine storage

- Enhanced networking support

For more information, see What's New in Hyper-V in Windows Server 2008 R2 (http://go.microsoft.com/fwlink/?LinkID=155234).

**See Also**

Virtualization support and licensing (SharePoint Server 2010)

# Plan virtual architectures (SharePoint Server 2010)

This article discusses key considerations for planning virtual architectures by using Microsoft SharePoint Server 2010 server roles. This article does not include performance or capacity planning data or recommendations. It describes general guidance for planning virtual environments and includes example architectures for small, medium, and large size farms.

In this article:

- Virtual versus physical architectures
- Example virtual architectures for small to medium size farms
- Example virtual architectures for medium to large size farms

## Virtual versus physical architectures

Typically, an organization considers a move to virtual architectures because it wants to reduce the number of servers that are required to host a solution, to more efficiently use existing hardware, or to save energy and space. The ability to automate server deployment is also a primary motivation to deploy a virtual server environment.

### Virtualizing Web servers and application servers

The Web server and application server roles are good candidates for virtualization. When you plan a virtual environment, a reasonable approach is to apply topology, performance, and capacity guidance to plan the physical environment, and then use the resulting number of Web servers and application servers — including specific application server roles — as a starting point for the virtual environment.

In a virtual environment; however, more virtual servers might be required to provide the same level of service and performance during peak times as are provided by physical servers. The results will depend on the specific services and the usage patterns of these services.

That said, running in a virtual environment provides the flexibility of re-allocating resources across virtual machines as necessary to tune performance. You can also more easily add and remove virtual servers to address spikes in usage of specific services that occur at predictable times throughout a year.

### Virtualizing SQL Server

The question of whether to virtualize Microsoft SQL Server is debatable and depends on the overall goals of a deployment. A virtual SQL Server environment typically performs a bit more slowly than a physical environment, although performance is improving as new versions are released. By using the most recent version of the Hyper-V role (included in Windows Server 2008 R2), SQL Server performance tests indicate that the same throughput (compared to a physical server) can be achieved in a guest virtual machine at a cost of slightly increased CPU usage.

There are other things to consider before you plan to virtualize SQL Server, such as the number of CPU cores required by SQL Server, the failover and availability plan, and the options for optimizing storage. Regardless, the benefits of deploying SQL Server to a virtual environment might outweigh the performance cost.

Organizations that host SharePoint farms and plan to deploy and rebuild farms often, such as hosting companies, will gain the most by adding SQL Server to the virtual environment. Virtualizing SQL Server might also be useful in a temporary or transitional solution, for example when combining multiple farms into an enterprise farm and retiring hardware. Organizations that make the most of limited hardware will gain the most by deploying SQL Server to physical servers. The examples in this article include environments that take both approaches.

For more information, see Running SQL Server 2008 in a Hyper-V Environment - Best Practices and Performance Recommendations (http://go.microsoft.com/fwlink/?LinkID=134106). This white paper is based on an earlier version of Hyper-V. Look for a newer version of this paper coming late Spring 2010.

## Virtualizing other servers in the environment

SharePoint 2010 Products solutions rely on other servers in the environment. This section provides general guidance on factoring these into a virtual architecture.

### Active Directory

We recommend that — at a minimum — the root domain controller for an Active Directory directory services environment be hosted on a physical server outside of virtual environments. If needed, additional domain controllers can be deployed as virtual servers.

For more information about how to deploy Active Directory to virtual environments, see the following resources:

- Blog by Sander Berkouwer: Active Directory in Hyper-V environments, Part 2 (http://go.microsoft.com/fwlink/?LinkId=186927)
- Planning Considerations for Virtualized Domain Controllers (http://go.microsoft.com/fwlink/?LinkId=186928)

### Gateway products

Gateway products include the following:

- Microsoft Forefront Unified Access Gateway (UAG)
- Microsoft Forefront Threat Management Gateway (TMG)

For high availability, we recommend that you position these products outside the SharePoint 2010 Products virtual environment. For more information about how to set up virtual environments for these gateway products, refer to the product documentation.

## Testing side by side

If you are concerned about how deploying SharePoint 2010 Products server roles in a virtual environment might affect performance, consider testing the specific roles that you plan to deploy. You can use the results to decide how many virtual servers to deploy for a specific role, or even whether to deploy a specific role to the virtual environment. For example, if your farm will crawl lots of content, the results of your tests might lead you to deploy the crawl role to a dedicated physical server.

One way to test a virtual environment is to deploy a specific role both virtually and physically, and compare benchmark data for network, memory, disk, and CPU. The following illustration provides an example of how to test specific server roles by using a limited number of servers.

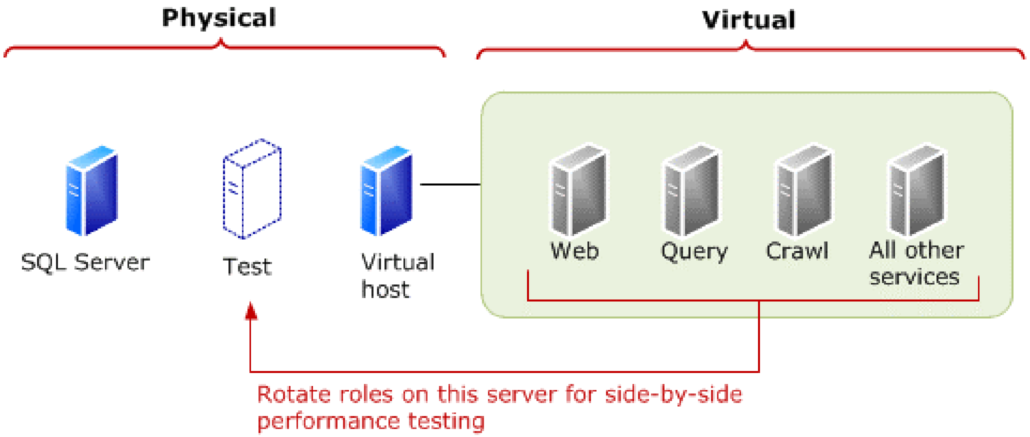

In this illustration, specific roles are deployed to the virtual environment. A physical test server is set up to test each role, one at a time, so that side-by-side benchmark data can be collected. Be sure to account for differences between the physical and virtual environments that will affect test results, such as different hardware specifications.

If you have an existing farm, you can add a virtual host and swap in virtual machines that have equivalent roles to see how the virtual performance for each role is affected. You can also see how different combinations of roles affect the overall performance of the farm. The following example illustrates this idea.

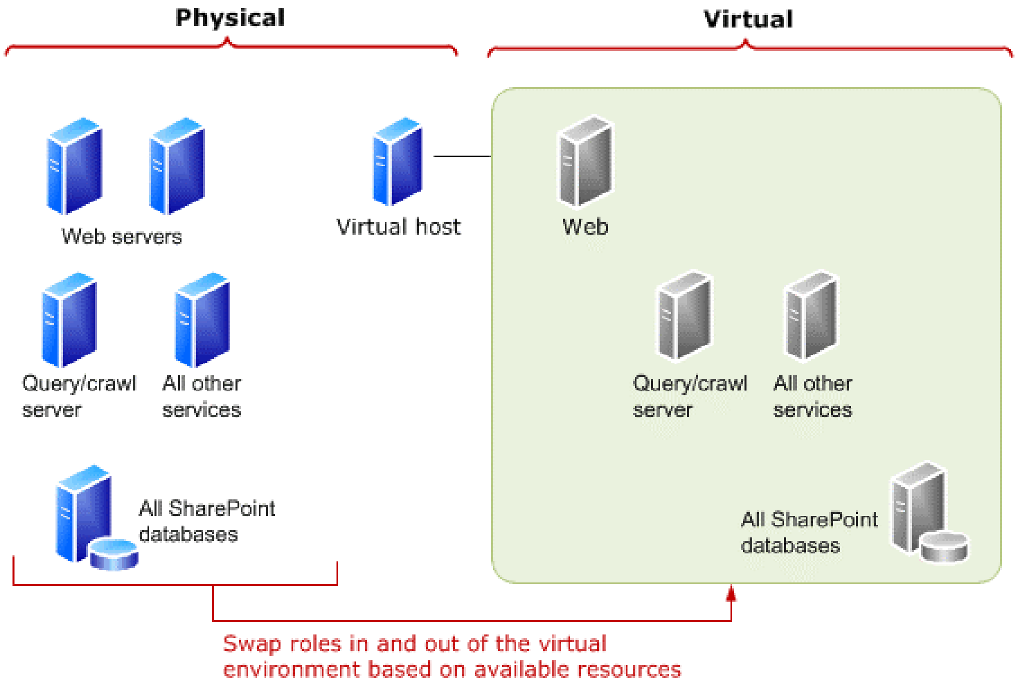

# Example virtual architectures for small to medium size farms

The starting point for replacing a physical farm by using a virtual farm is to use two to four physical host servers. For each host, the number of servers that can be deployed is dictated by available memory, CPU, disk, and network resources.

The following two illustrations provide example deployments in which the Web server and application server roles are deployed to a virtual environment.

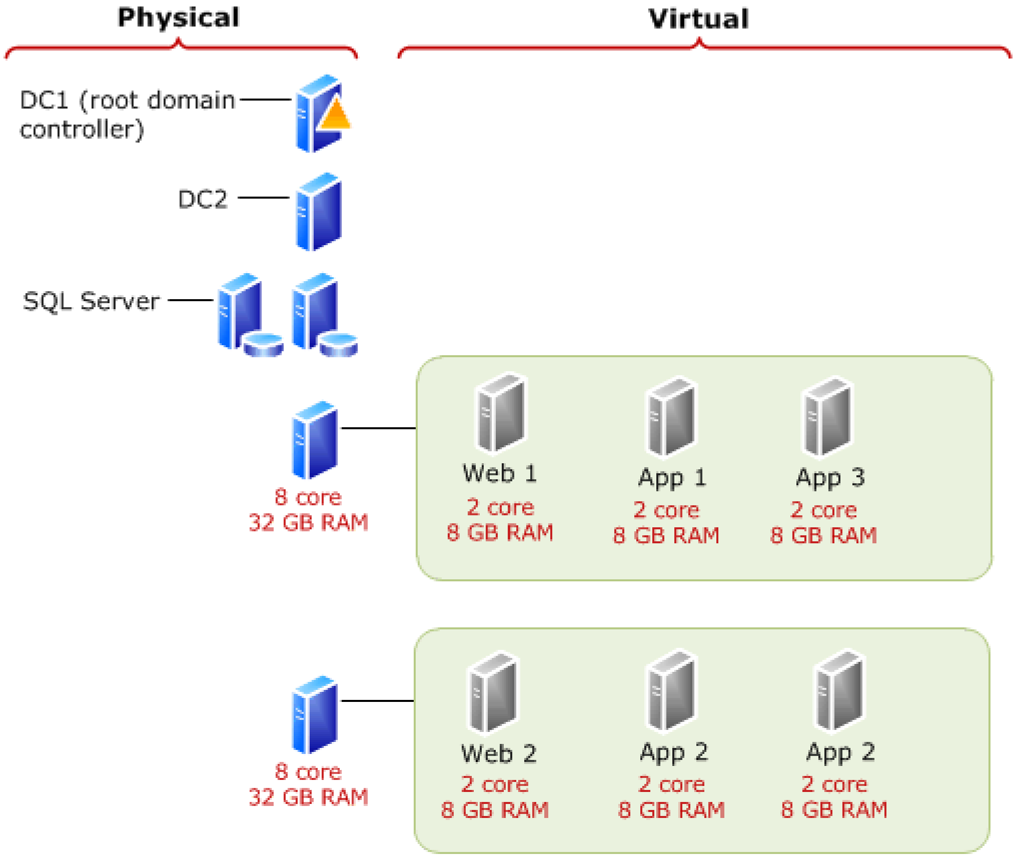

In this example, be aware of the following:

- The minimum resources for CPUs and RAM represent starting points for the farm. Because only two cores are reserved for each virtual image, this example is only appropriate for proof-of-concept or development environments in which performance is not an issue. Reserve enough spare resources to reallocate based on performance monitoring.
- SQL Server is deployed to physical servers, instead of to virtual servers.
- Web servers and application servers are redundant across the two host servers.
- Three Web servers are deployed to the virtual environment for high availability.
- The Active Directory domain controllers are deployed to physical servers.

For pilot testing and production environments, four cores is the minimum recommended starting point for virtual machines. The following virtual environment uses fewer virtual machines to achieve this objective.

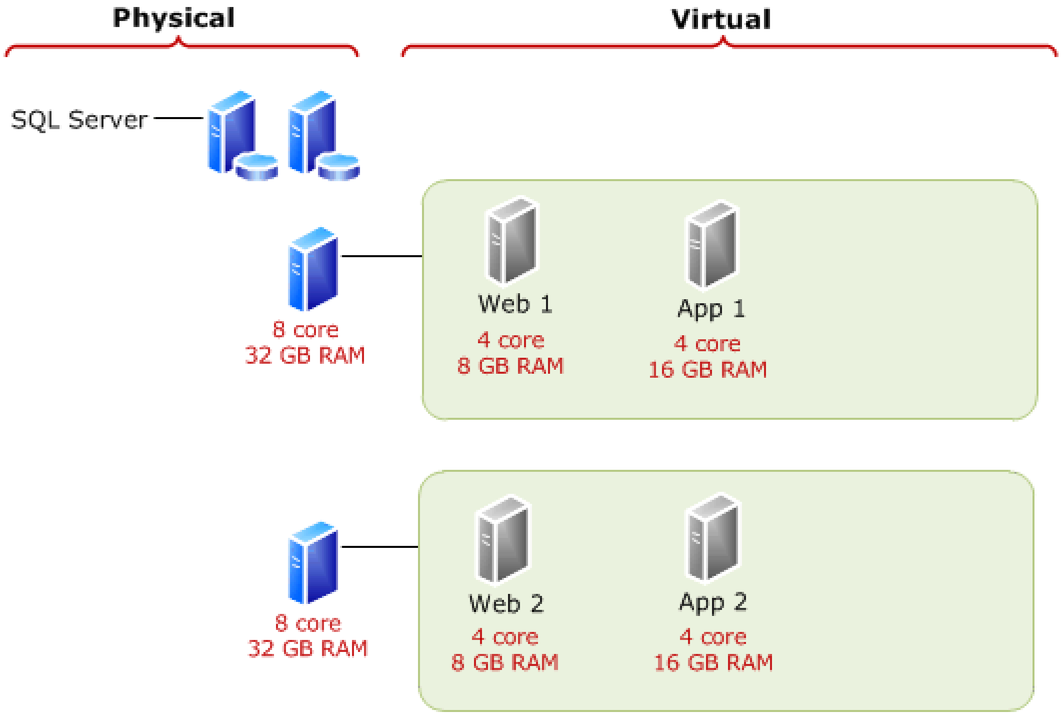

This example represents a starting-point environment. You might have to add resources, depending on the usage pattern of the farm.

## Example virtual architectures for medium to large size farms

By using larger host servers, you can allocate more resources to virtual images. The following illustration provides an example implementation that uses more CPUs and RAM.

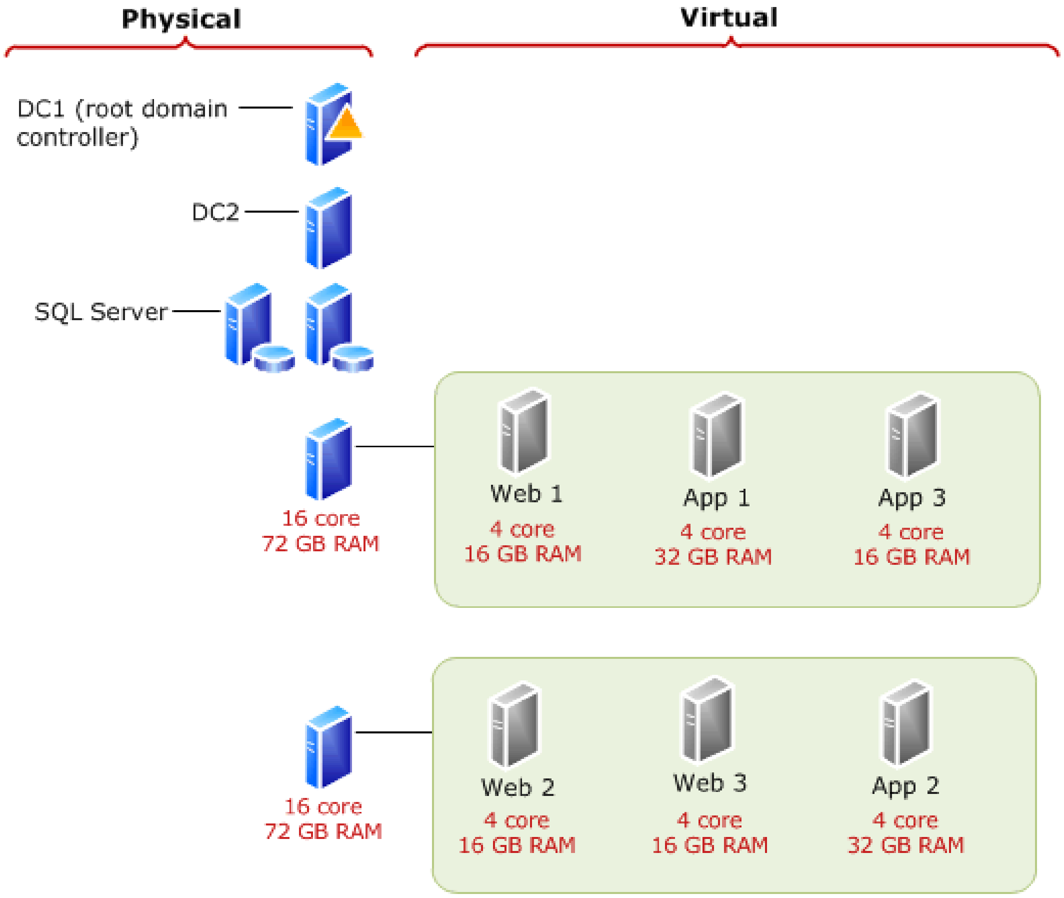

If the benefits of virtualizing SQL Server outweigh the performance tradeoffs, SQL Server can be deployed as a guest also, as shown in the following illustration.

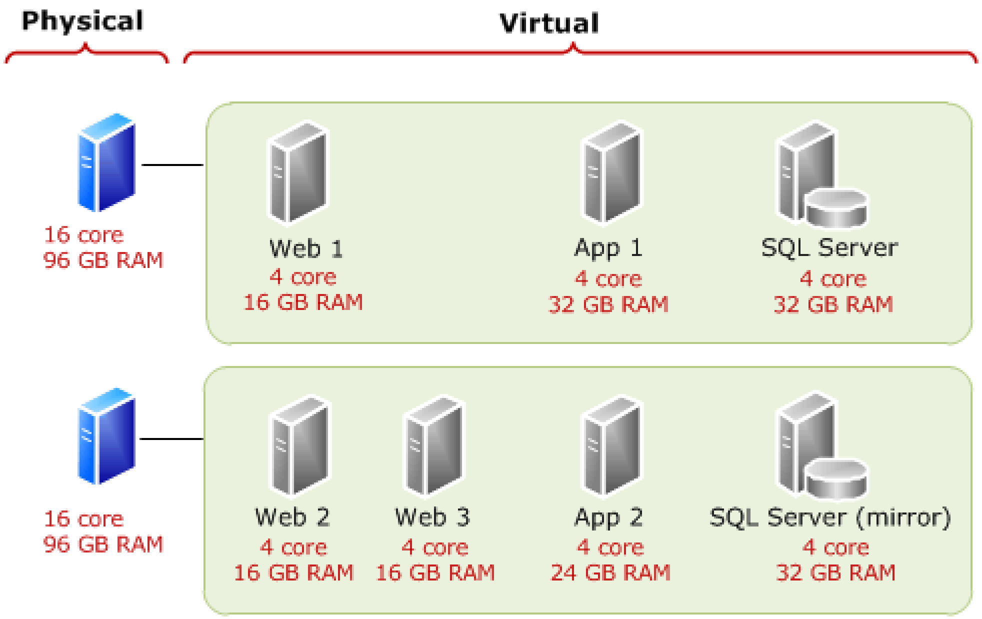

In this example, be aware of the following:

- Only one instance of SQL Server is deployed to each host. In small and medium size virtual environments, we recommend that you do not deploy more than one SQL Server guest per host.
- Both host servers include more memory to accommodate the number of virtual servers, including SQL Server.

If a particular server role consumes so many resources that it adversely affects the overall performance of the virtual environment, consider dedicating a physical server to this role. Depending on the usage patterns of an organization, these roles might include crawl servers, the server that imports profiles, Excel Services Application, or other services that are used heavily. The following illustration provides an example.

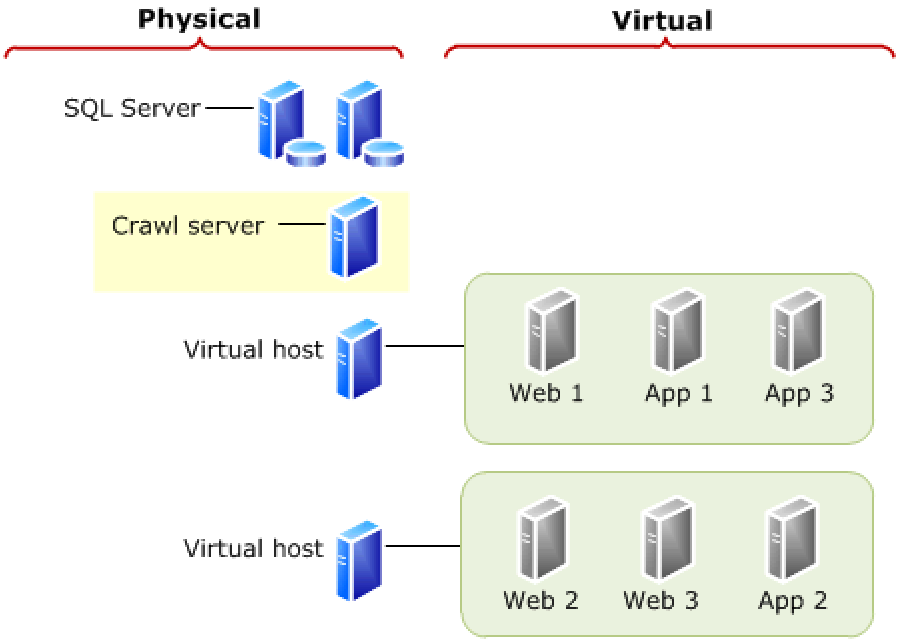

In this example:

- SQL Server is deployed to physical servers. Remove SQL Server from the virtual environment before you remove application server roles.
- The crawl role is deployed to a physical server. In some environments, a different role might be a candidate for deploying to a physical server, depending on usage.

# Plan for virtualization (SharePoint Server 2010)

This article describes the planning process to follow in order to successfully deploy Microsoft SharePoint Server 2010 in a virtual environment. Each step in the planning process includes links to the appropriate documentation. It is assumed that you have determined the SharePoint Server 2010 solution that you want to deploy in a virtual environment. On the surface, deploying a SharePoint Server 2010 farm on virtual machines is the same as deploying a farm on physical servers. However, deploying in a virtual environment involves a different level of planning that takes into account the characteristics of Windows Server 2008 Hyper-V technology as well as how virtual machines, the virtual network adapters, and virtual hard disks are implemented on a virtualization server.

Before you start developing your virtualization plan, we recommend that you read the Hyper-V Planning and Deployment Guide (http://go.microsoft.com/fwlink/?LinkId=187964).

Detailed information about the following subjects is out-of-scope for this article, but is provided in other articles:

- Capacity management
- Security requirements
- Health and performance monitoring
- Backup and recovery

A virtual environment consists of two interrelated layers, one physical and one virtual. A configuration change in either layer effects servers in the other layer. This interrelationship becomes evident when you plan for, deploy, and use SharePoint Server 2010 in a virtual environment. For more information, see Virtualization overview.

# Create a plan for deploying SharePoint Server 2010 in a virtual environment

You should approach planning for a virtual farm the same way as you would plan for a physical farm. Most, if not all, of the issues and requirements for deploying SharePoint Server 2010 on physical servers apply equally to virtual machines. Any decisions that you make, such as minimum processor or memory requirements, have a direct bearing on number of virtualization hosts required, as well as their ability to adequately support the virtual machines that you identify for the farm.

After you finish planning a physical farm, you have all the information you need to design virtualization architecture. Ideally, this architecture is as close as possible to the final virtualization solution that you intend to put into production. Realistically, the architecture is likely to change as you move through the deployment phase of the system lifecycle. In fact, you may determine that some farm server roles are not good candidates for virtualization.

The key planning steps, tasks, and references are summarized in following the procedure.

### To create a virtualization plan

1. Determine virtualization scope

   Determining the scope of farm virtualization is a key contributing factor to successfully implementing, managing, and evaluating your virtualization project. When determining scope, you have to decide whether you will virtualize some or all of the supporting virtual machine infrastructure.

   Use the following list of tasks to determine the scope of virtualization.

   - Task 1: Identify all the farms that are required to implement your solution. Take into consideration the fact that most solutions have several farm components. For example, an Internet-facing Web portal typically has a publishing farm, an authoring farm and a testing or quality assurance farm.
   - Task 2: For each farm, determine the number of servers that are required as well as the role that each server will have in the farm.
   - Task 3: Identify which farms you want to deploy in in a virtual environment.

   Refining the scope of a solution also refines the scope of a deployment, which makes it easier to implement and manage. For more information, see Plan for sites and solutions (SharePoint Server 2010). In many cases, solutions share common elements; however, each solution may have its own requirements. For more information, see Fundamental site planning (SharePoint Server 2010). The Plan for social computing and collaboration (SharePoint Server 2010) article shows one of the popular solutions.

   **Note:**
   Expect to refine the scope of your solution as you move through the phases of deploying your farm in a production environment.

2. Identify servers to virtualize

   Identify servers that are good candidates for virtualization. From a technical and Microsoft support perspective, all SharePoint servers can be virtualized. The decision to virtualize a particular farm server should be based on:

   - Corporate compliance policies (for example, legal and technical)
   - Benefits derived from server consolidation, such as reduced power consumption and physical space requirements. For more information, see Server virtualization (http://go.microsoft.com/fwlink/?LinkId=187965).
   - Capacity requirements (see next planning step)

3. Identify capacity requirements for each farm server

   Determine the resource requirements for each farm server as if it was a physical server. Take into account specialized server roles, such as hosting Enterprise Search components. You need to specify the amount of resources needed for each of the following server components:

   - Memory

- Number of processors and minimum clock speed
- Number and size of hard disks
- Number of network adapters and their required throughput speed

Use the resources available at the Capacity Management for SharePoint Server 2010 Resource Center to develop the capacity and performance requirements for your farm.

4. Determine if virtual machine can meet physical requirements.

   You have to determine whether each virtual machine that you identified in Step 3 can meet the capacity requirements of a corresponding physical server. At a minimum, complete the following tasks:

   - Task 1: Assess the memory requirement in the context of available virtualization host capacity.
   - Task 2: Assess the processor requirement. Hyper-V has a hard limit of four virtual processors per virtual machine. If a physical farm server requires eight processors, determine whether this requirement can be met by scaling out the number of virtual machines in a farm.
   - Task 3: Assess the virtual machine storage requirement in the context of local physical storage or SAN.

5. Determine virtualization host requirements

   Determine the minimum host requirements (memory, number of cores, number and size of local hard drives, number of network adapters)Also consider and plan for the following:

   - Scalability: Determine if you can add more CPUs, more memory, more hard disks, and more network adapters to the host computer.

     **Important:**
     Depending on the manufacturer and computer model, you may not be able to increase capacity. You need to have this information before you use or purchase a server.

   - Extra host capacity: Determine whether or not the host has the capacity to scale up existing virtual machines, or to add additional virtual machines. This is very important if you plan to use Hyper-V failover clustering, quick migration, or live migration.

   **Important:**
   Plan for peak load and determine how short term spikes in load will be handled.

6. Design virtualization architecture

   A well-designed architecture is required for a successful solution. For SharePoint Server 2010, a basic three-tier topology provides the foundation for all the solutions. The following elements form a good design that is based on the recommended foundation topology:

   - Good overall performance
   - Ease of maintenance and upgrade

   - Flexibility
   - Scalability
   - High availability

   For more information, see Plan for server farms and environments (SharePoint Server 2010).

   A virtualization architecture model consists of the virtualization hosts and the virtual machines that make up the farm topology. This model enables you to visualize the virtual environment that you plan to deploy. For more information about virtualization architectures, see Plan virtual architectures (SharePoint Server 2010) and Capacity management and high availability in a virtual environment (SharePoint Server 2010).

   **Note:**
   Be prepared to refine the architecture as you move through the planning process. The following steps may dictate changes to the architecture.

7. Identify storage requirements

   Determine how much local physical storage or SAN storage is required for Hyper-V-related storage such as configuration files, Virtual Hard Disks (VHDs), and snapshots.

8. Identify backup and recovery requirements

   In addition to the farm servers, you have to plan backup and recovery for all or part of a farm. For more information, see Backup and recovery (SharePoint Server 2010).

9. Determine high availability requirements and design a solution

   Identify approaches for achieving high availability for Web servers, application servers, and databases. Typical strategies include the following:

   - Redundant hardware and servers
   - Hot-swappable components
   - Failover clustering for virtual and physical servers. For more information, see Hyper-V: Using Hyper-V and Failover Clustering (http://go.microsoft.com/fwlink/?LinkId=187967).
   - Clustering or mirroring for database servers. For more information, see Plan for availability (SharePoint Server 2010)

10. Identify health and capacity indicators for monitoring the virtual environment. For more information, see Capacity management and high availability in a virtual environment (SharePoint Server 2010).

    Combine the key indicators that you derived in the previous steps with the planning you did for SharePoint Server 2010. For more information, see Plan for server farms and environments (SharePoint Server 2010). You have to determine all the health and capacity indicators in order to collect measurements from the following objects in the virtual environment:

    - Virtual machines with SharePoint Server 2010 installed
    - Virtual machines that are not part of the farm, such as a firewall server
    - Virtualization hosts

   - Network components

   After you start to collect data from the virtual environment, you can create a baseline, which can be used to assess and tune the virtual environment during deployment and after the farm goes into production.

11. Create a deployment plan for the deployment phase of the system lifecycle.

    For more information, see the SharePoint 2010 Products Deployment model, available in the Technical diagrams (SharePoint Server 2010) article.

12. Create a maintenance plan

    Create a maintenance plan that enables you to implement password changes and apply software updates, service packs, and hotfixes. This plan should include the virtual machines and the virtualization hosts.

# Capacity management and high availability in a virtual environment (SharePoint Server 2010)

This article provides information about capacity management and high availability for a virtual environment hosting Microsoft SharePoint Server 2010. We combine these two concepts in this article because capacity and sizing are very important parts of developing a virtualization plan and the architecture for a virtual environment, and because capacity management is not isolated from high availability in a virtual environment. In the case of virtualization hosts, insufficient capacity can block high availability at the farm level and at the host level.

As is the case with other aspects of a virtual environment, such as backup and recovery, capacity management and high availability have to accommodate the two layers of a virtual environment, which are the virtual machines used for SharePoint Server 2010 and the physical servers that are used to host the virtual machines. In the case of a hybrid environment, you also have to deal with physical Microsoft SharePoint Server 2010 farm servers.

In this article:

- Virtualization overview
- Capacity management
- Virtualization server capacity and sizing
- Creating and refining the architectures
- Additional options for improving the architecture

## Virtualization overview

Server virtualization, as implemented by Windows Server 2008 Hyper-V technology or Microsoft Hyper-V Server 2008, is hardware-based and is also referred to as hardware-assisted virtualization, as opposed to software-based virtualization. The Hyper-V hypervisor has a more direct communication path to, and interaction with, physical server hardware components than software-based virtualization technologies. The net result is better performance than a software-based virtualization technology. For more information about the Hyper-V architecture, see An Introduction to Hyper-V in Windows Server 2008 (http://go.microsoft.com/fwlink/?LinkId=188006) and Monitoring Hyper-V performance (http://go.microsoft.com/fwlink/?LinkID=187746&clcid=0x409).

Although a physical server may meet Hyper-V requirements, each physical server is unique. Every manufacturer uses its own implementation of processors, multi-core technology, memory, the data bus, hard disks, and network adapters. Additionally, hardware design and implementation varies from model to model, even if the models are produced by the same manufacturer. This highlights the need for rigorous testing when you deploy SharePoint Server 2010 in a virtual environment.

Software programs and applications exhibit the same variations in performance as hardware. Some programs are CPU intensive, other programs have high memory demands, and other programs are hard disk intensive. SharePoint Server 2010 has its own capacity needs, as does Internet Information Server (IIS) and SQL Server 2008. Once again, rigorous testing is required.

Capacity management requires you to consider the virtualization server, the storage solution, the network infrastructure, the technologies running in a SharePoint Server 2010 environment, and the features that are enabled to implement your SharePoint Server 2010 solution.

# Capacity management

Capacity management extends the concept of capacity planning to express a cyclical approach in which the capacity of a SharePoint Server 2010 deployment is continually monitored and optimized to accommodate changing conditions and requirements. You can apply this approach to all SharePoint Server 2010 farms, including those that are fully virtualized, and those that are partly virtualized. For an overview of capacity management, see Capacity management and sizing for SharePoint Server 2010. Additional capacity management resources are located at the Capacity Management for SharePoint Server 2010 (http://go.microsoft.com/fwlink/?LinkId=194748) Resource Center.

# Virtualization server capacity and sizing

After you have a SharePoint Server 2010 farm design and sizing recommendations for the farm servers, you design the physical virtualization host architecture required to support the virtual farm. For more information about virtual architectures, see Plan virtual architectures (SharePoint Server 2010).

We recommend that you use applicable principles from SharePoint Server 2010 capacity management and use them as guides for a virtual environment. The following activities illustrate the iterative nature of designing, sizing, and adjusting a virtual and physical architecture from initial planning to deployment in a production environment.

**Note:**
If you do thorough planning and testing, changes to the architecture and server configurations should only be required if there is a significant and unanticipated increase in farm use or because new features are added to your SharePoint Server 2010 solution.

- Before you start the farm deployment, create a virtual and physical architecture with virtual machine and virtualization server sizing. In the case of multiple virtualization hosts, this architecture must include virtual machine distribution.
- During the pilot phase of deployment, collect health and performance data that you can use to establish benchmarks for the farm virtual machines and the virtualization hosts.
- During the user acceptance test phase of deployment, adjust the virtualization host and virtual machine configurations based on the benchmark data. If necessary, change the physical architecture by redistributing virtual machines on the virtualization hosts.
- After deployment, continue to collect health and performance benchmarks and refine virtual machine and, if applicable, physical machine configurations. If necessary, adjust both architectures.

It is essential that you can analyze virtualization host and virtual machine performance data and understand how it reflects capacity needs and application effect on capacity. Additionally, you have to understand performance and capacity limits. Given the interrelationship between the virtual layer and the physical layer, anything that affects virtual machine capacity and performance either has a direct effect on the host or has to be accommodated for by making changes to the virtualization host configuration in order to sustain acceptable performance across the farm.

In some cases it may be necessary to change the physical architecture by adding additional virtualization hosts and then changing the distribution of virtual machines on the physical architecture.

**Important:**

In benchmark tests between a physical computer and a virtual machine, the throughput of the virtual machine typically cannot match that of a physical computer. Virtual machine performance is, with rare exceptions, always lower than that of a physical computer. The degree of performance difference depends on virtualization host capabilities, the applications that are running, and the benchmarks that you choose to use as primary performance indicators.

We recommend that you read the Hyper-V Performance FAQ R2 (http://go.microsoft.com/fwlink/?LinkID=187745), which is updated to reflect capacity and performance information for Windows Server 2008 R2 and Windows Server 2008 with Service Pack 2 (SP2). This FAQ contains answers to common Hyper-V questions, provides guidance, and includes links to detailed articles that you can use to develop benchmarks for the virtualization host, virtual machines, and Windows networking.

We also suggest you read the following posts about Hyper-V performance counters:

- Hyper-V Performance Counters - Part one of many - The overview (http://go.microsoft.com/fwlink/?LinkID=125651)
- Hyper-V Performance Counters – Part two of many – "Hyper-V Hypervisor" counter set (http://go.microsoft.com/fwlink/?LinkID=125652)
- Hyper-V Performance Counters – Part three of many – "Hyper-V Hypervisor Logical Processors" counter set (http://go.microsoft.com/fwlink/?LinkID=125653)
- Hyper-V Performance Counters – Part four of many – "Hyper-V Hypervisor Virtual Processor" and "Hyper-V Hypervisor Root Virtual Processor" counter set (http://go.microsoft.com/fwlink/?LinkID=125655)

# Creating and refining the architectures

A complete architecture consists of the virtualization hosts, virtual machines, and physical machines that make up the SharePoint Server 2010 environment that you plan to deploy. For more information about virtualization architectures, see Plan virtual architectures (SharePoint Server 2010).

Developing and implementing a virtual architecture consists of the following steps:

1. Create the virtual and physical architecture. Create an architecture that will support the goals of your SharePoint Server 2010 farm.

2. Analyze the architectures. Identify and obtain any information that is missing or that will improve the design of the environment that you plan to deploy.
3. Refine the architectures. Use the information from step 2 to refine the architecture.
4. Continue to refine the architectures and server configurations as you move through the various deployment stages. For more information about the deployment stages, see the SharePoint 2010 Products Deployment and SharePoint 2010 Products: Virtualization Process models, available in the Technical diagrams (SharePoint Server 2010) article.

## Create the architecture

Create a model of the architecture that you can use as a tool for evaluating and adjusting virtual machine and virtualization host configurations. Use the following criteria as a guide for developing your model:

- Identify the number of virtual machines that are needed and the role of each in the SharePoint Server 2010 farm.
- Specify individual virtual machine configuration requirements (disk space, memory, and number of processors). These are based on SharePoint Server 2010 capacity requirements.
- Specify virtualization host requirements (disk space, memory, and number of logical processors). These are based on virtual machine requirements.
- Identify virtual machine distribution on virtualization hosts. These are based on farm high-availability requirements and are constrained by virtualization host quantity and capacity.
- Identify general networking and storage requirements.
- Allow for growth on virtualization hosts and on the virtual machines (scale up or scale out).

After you create an architecture model, you have to analyze both architectures to validate the design as well as the virtualization host and virtual machine configurations.

## Analyze the architectures

The fundamental purpose of analyzing the architecture is to determine whether it can successfully support the SharePoint Server 2010 solution that you want to deploy. However, it is reasonable to assume that design and server configurations will change as you move through the deployment process.

The following illustration shows a sample virtual architecture for a farm that consists of front-end Web servers, application servers, and database servers. This architecture is representative of the small to medium farms described in Example virtual architectures for small to medium size farms, and we can use it show the key elements that have to be considered when you analyze the capacity and availability requirements for a virtual farm.

**Important:**
Virtualization server and virtual machine sizing in the following illustration is not prescriptive.

**Figure 1. Preliminary architecture**

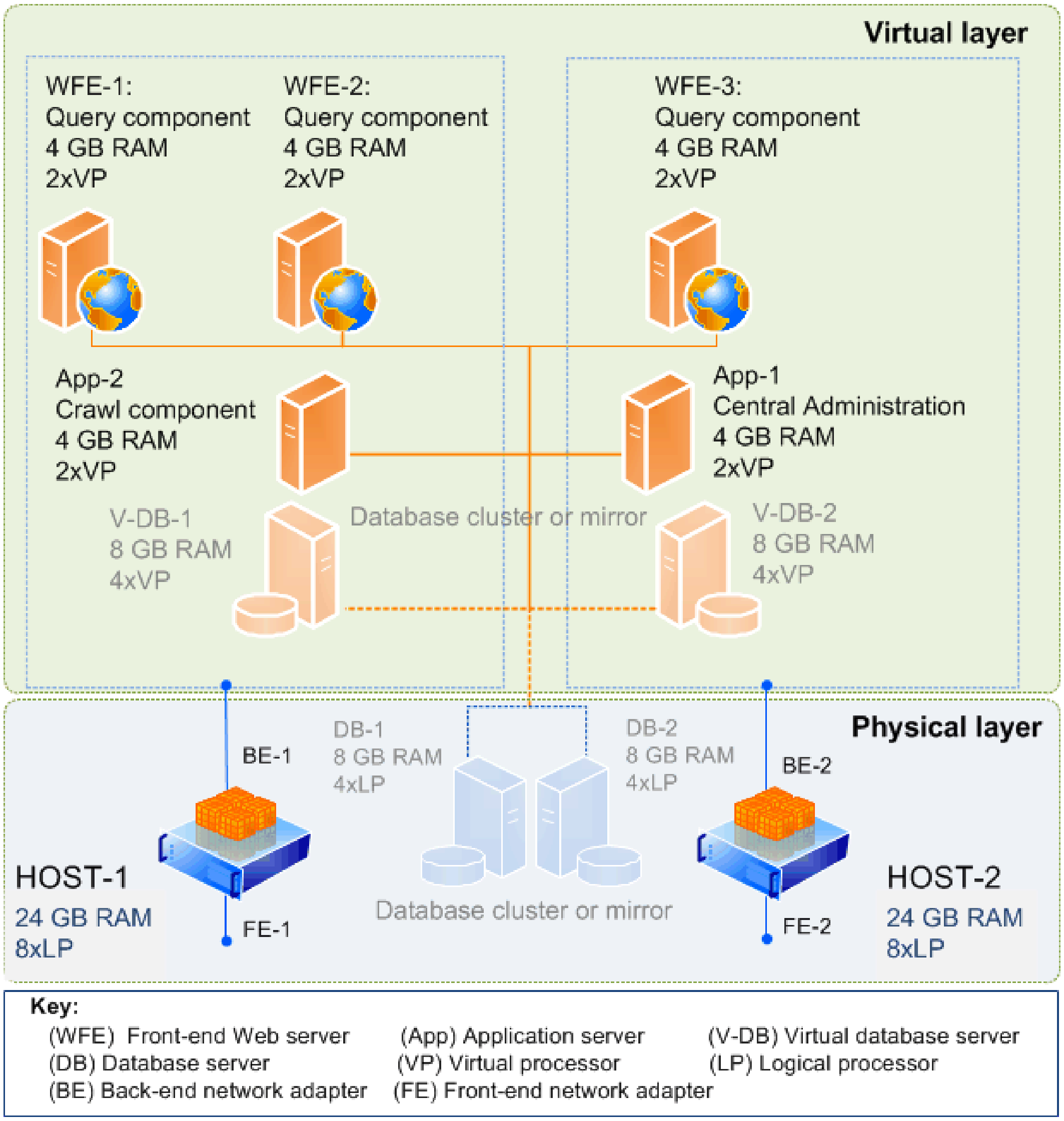

Use the criteria provided for creating a virtual architecture to analyze the sample architecture shown in the previous illustration. The architecture in the illustration assumes that all the Web servers and application servers are virtual machines. It has not been determined whether the farm database servers are physical machines or virtual machines.

### Virtualization host analysis

The following tables (HOST-1 and HOST-2) provide an analysis of each virtualization host and use memory, processors, and scalability as criteria. The host analysis is followed by a design analysis.

**HOST-1**

| Criteria | Analysis |
|---|---|
| Memory | After factoring in 2 GB of RAM for the host operating system and using the projected RAM requirements, there is an estimated 2 GB of RAM available for future use. |
| Processors | The logical to virtual processor mapping is 8:10 (1:1.25), which means that the CPU is slightly oversubscribed, which would not be an issue in a test environment.<br>**Important:**<br>Oversubscribing the CPU on a virtualization server will reduce overall performance. The extent of this effect is determined by the load put on the virtual machines. As a best practice, do not oversubscribe the virtualization server CPU if it can be avoided. |
| Scalability | This is not an option because there is insufficient memory. Additionally, the degree of CPU oversubscription (even by adding a virtual machine with two processors) would have a noticeable effect on performance. |

**HOST-2**

| Criteria | Analysis |
|---|---|
| Memory | After factoring in 2 GB of RAM for the host operating system and using the projected RAM requirements, there is an estimated 6 GB of RAM available for future use. |
| Processors | The logical to virtual processor mapping is 8:8 (1:1), which meets the best practice guidance. |
| Scalability | There is sufficient memory to increase the memory allocation to the virtual machines. There is enough capacity to add a new virtual machine with two processors and 4 GB of RAM. This means that the |

| Criteria | Analysis |
|---|---|
| | virtualization host CPU would be slightly oversubscribed (8:10), but like HOST-1, would not be an issue in a test environment. |

**Design analysis**

The sample architecture generally shows a degree of high availability for the farm servers. For example, there are three front-end Web servers distributed across HOST-1 and HOST-2, and the database servers (clustered or mirrored) also reside on separate virtualization hosts or separate physical servers. High availability at the virtualization host level is not part of the architecture and pertinent information is missing. The following information is required before the design can be revised:

- Database size

  The size of the content database determines how you configure and distribute all the farm servers.

- Storage subsystem

  For example, in the sample architecture, no information is provided about the number of disks required for each virtual machine, nor is there any indication of disk distribution and capacity. This information is very important for determining and configuring the storage system. The architecture sample uses local storage. You have to determine whether this is suitable for your environment, or if you want to use pass-through disk configuration to a LUN on a SAN.

- Networking requirements

  The number of network adapters and minimum throughput has to be identified.

- Virtual hard disk configurations

  You also have to determine which of the Hyper-V hard disk configurations you want to use (for example, fixed size, pass-through). For more information, see Planning for Disks and Storage (http://go.microsoft.com/fwlink/?LinkId=188007) and Virtual Hard Disk Performance: Windows Server 2008 / Windows Server 2008 R2 / Windows 7 (http://go.microsoft.com/fwlink/?LinkId=186519).

After you complete a design review, the next step is to refine the architecture.

## Refine the architecture

The scope of refining the architecture depends on your initial architecture, the results of your analysis, and your implementation plan. Using the sample that is provided, there are scenarios where you might decide not to make any changes. For example:

- The preliminary architecture is suitable for early testing, proof-of-concept, and a limited pilot deployment.
- The virtualization hosts are for testing only and will be replaced by higher capacity hosts during the user acceptance test phase.

- The virtual farm is for testing purposes only and will be shut down after testing is finished. In some cases the environment may be retained and used at a later date to test software updates.

The following illustration shows a revised architecture that is better suited for a production farm.

**Figure 2. Revised architecture**

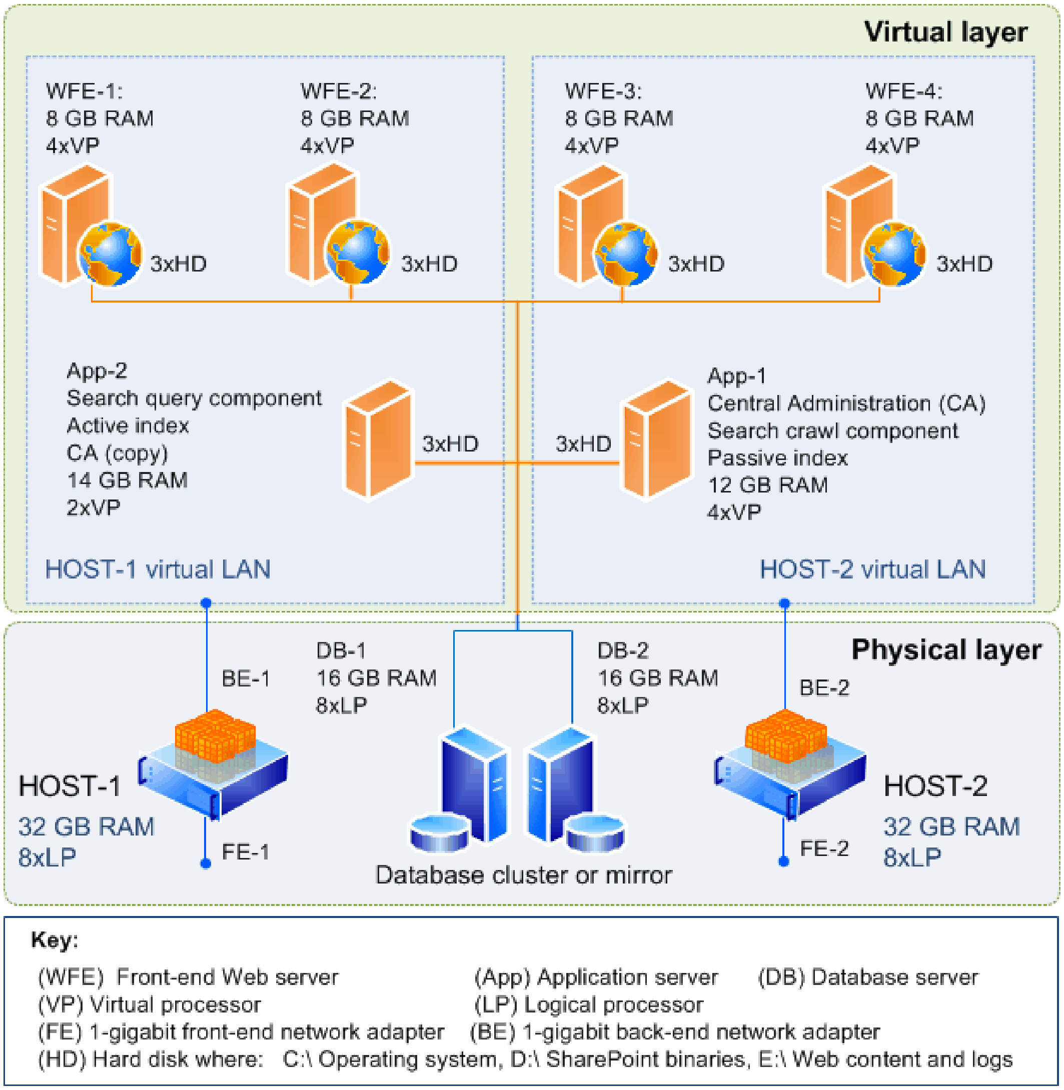

In the revised architecture, the primary assumption is that you want to stay with eight core commodity virtualization servers. The changes in the preceding illustration reflect that assumption and include the following considerations:

- The estimated size of the content database is 1 terabyte (TB).
- The objective is to provide high availability for all the farm servers and to maximize performance across infrastructure.

- The farm database servers are physical servers that can be clustered or mirrored to support high availability. Each server has 8 cores, 16 GB of RAM and uses local drives to reduce latency.

### Virtualization host analysis

The following tables (HOST-1 revised and HOST-2 revised) provide an analysis of each virtualization host using memory, processors, and scalability as criteria. The host analysis is followed by a design analysis.

**HOST-1 revised**

| Criteria | Analysis |
|---|---|
| Memory | After factoring in 2 GB of RAM for the host operating system and using the projected RAM requirements, there is an estimated 2 GB of RAM available for future use. |
| Processors | The logical to virtual processor mapping is 8:10 (1:1.25), which is slightly oversubscribed. |
| Scalability | There is a marginal amount of memory available to increase the memory allocation to the virtual machines. Based on the amount of memory and the processor ratio, there is not enough host capacity to add an additional virtual machine. |

**HOST-2 revised**

| Criteria | Analysis |
|---|---|
| Memory | After factoring in 2 GB of RAM for the host operating system and using the projected RAM requirements, there is an estimated 4 GB of RAM available for future use. |
| Processors | The logical to virtual processor mapping is 8:12 (1:1.50), which is oversubscribed by 50 percent. |
| Scalability | There is a marginal amount of memory available to increase the memory allocation to the virtual machines. Based on the amount of memory and the processor ratio, there is not enough host capacity to add an additional virtual machine. |

### Design analysis

- Each virtual machine uses a three-drive configuration, sized according to SharePoint Server 2010 best practice guidance. These drives are typically configured as follows:
  - Drive C  (50 GB) for Windows installation
  - Drive D (50 GB) for SharePoint Server 2010 files
  - Drive E (300 GB) for Web content and log files
- Each front-end Web server is configured with four virtual processors (4xVP) and 8 GB of RAM. This is the minimum recommended configuration for a production environment.
- The number of front-end Web servers is increased to four to support effective clustering and high availability. This four-server configuration is particularly well-suited for installing software updates because there will always be two servers available when you install updates.
- The two application servers (App-1, App-2) provide high availability. App-1 hosts Central Administration, the Search crawl component, and the passive index for the Search query component. The number of processors and amount of memory is based on the estimated size of the content database.

  App-2 is a dedicated Search query server. It also contains a copy of Central Administration. The number of processors and amount of memory is based on the estimated size of content database.
- For high availability, Central Administration is also installed on a front-end Web server on another host.
- The database servers are physical servers that are clustered or mirrored to ensure high availability. This move to physical servers has the benefits of increasing virtualization host capacity for the virtual farm servers and improving overall database performance.

  **Note:**
  As indicated earlier in this article, the decision to virtualize or not virtualize database servers is a complex decision that requires extensive planning and testing.
- From a networking perspective, both virtualization hosts are configured with two separate 1-gigabit physical network adapters. This is a recommended practice to ensure that virtualization host and virtual machine data traffic is separated to improve performance and provide some adapter redundancy.
- Each virtualization host employs a virtual LAN (VLAN), which can provide the following benefits: network segregation, improved security, and performance.

The revised virtual and physical architecture is significantly improved and could be deployed into a production environment. However, it is important to note that, as configured, available virtualization host resources do not support farm scaling. Additionally, they cannot support the migration of a farm server from one host to another if the need arises.

Realistically, if you want to deploy the example farm into production, we recommend that you consider the following upgrades:

- Increase virtualization host capacity by using a 16-core computer with 48 or 64 gigabytes of RAM.
- Add one or more virtualization hosts.

To achieve the optimum level of high availability, consider the additional options in the following section.

# Additional options for improving the architecture

The previous section provided options for revising the model. There are, of course, other options to achieve better performance and high availability. Scaling out the virtualization host environment or scaling up the virtualization hosts are good alternatives, although cost is always an issue. Your organization's virtualization strategy will help define the best approach.

**Tip:**
In terms of cost, it is usually less expensive to purchase a server that has more capacity than you need in the short term than it is to upgrade a server to gain more capacity. This is especially true in the case of memory upgrades, where typically you have to throw away the existing memory modules and buy a full set of new memory in order to upgrade the memory.

Performance gains can be achieved with the following options:

- Deploy or purchase servers that have Second-Level Address Translation (SLAT) enabled processors. In Intel processors, this feature is referred to as Nested Page Tables and is available in Nehalem 55xx series processors. For AMD, this feature is referred to as Enhanced Page Tables (EPT).
- Deploy or purchase servers that provide CPU Core-Parking, a feature that allows the running Hyper-V to use the lowest number of processor cores to meet the workload demand.
- Investigate TCP chimney offload, Virtual Machine Queues (VMQ) and jumbo frames. These features improve network performance and decreases CPU utilization, thereby increasing the overall system capacity.
- Investigate jumbo frame support to speed up network performance when transferring large amounts of data. However, you must test this thoroughly because jumbo frames do not work in all environments.
- Investigate adapter teaming. This feature can improve network performance and provide failover capability to the physical network adapters.

  **Important:**
  Adapter teaming is a third-party solution and is only supported by the vendor. For more information, see Microsoft Support Policy for NIC Teaming with Hyper-V (http://go.microsoft.com/fwlink/?LinkId=194749).

To ensure high availability for a virtual environment, consider implementing Windows Server 2008 R2 failover clustering and Hyper-V live migration, as follows:

- The scope of failover clustering can include virtualization hosts and the virtual machines on each host. If a virtualization host fails unexpectedly, the virtual machines automatically fail over to another virtualization host.
- Live migration is a solution for planned downtime. You can migrate running virtual machines to another server (without downtime), shut down the physical server, and perform the maintenance.

When you finish maintenance on the server, use live migration to move the virtual machines back to the original physical server.

For more information, see Hyper-V: Using Hyper-V and Failover Clustering (http://go.microsoft.com/fwlink/?LinkID=187967&clcid=0x409) and Hyper-V: Using Live Migration with Cluster Shared Volumes in Windows Server 2008 R2 (http://go.microsoft.com/fwlink/?LinkId=188009).

# Plan authentication (SharePoint Server 2010)

In this section:

- Plan authentication methods (SharePoint Server 2010)
- Plan the Secure Store Service (SharePoint Server 2010)

# Plan authentication methods (SharePoint Server 2010)

This article describes the authentication methods and authentication modes that are supported by Microsoft SharePoint Server 2010. Authentication is the process of validating a user's identity. After a user's identity is validated, the authorization process determines which sites, content, and other features the user can access. Authentication modes determine how accounts are used internally by SharePoint Server 2010.

In this article:

- Supported authentication methods
- Authentication modes — classic or claims-based
- Implementing Windows authentication
- Implementing forms-based authentication
- Implementing SAML token-based authentication
- Choosing authentication for LDAP environments
- Planning zones for Web applications
- Architecture for SAML token-based providers

## Supported authentication methods

SharePoint Server 2010 supports authentication methods that were included in previous versions and also introduces token-based authentication that is based on Security Assertion Markup Language (SAML) as an option. The following table lists the supported authentication methods.

| Method | Examples | Notes |
|---|---|---|
| Windows | • NTLM<br>• Kerberos<br>• Anonymous<br>• Basic<br>• Digest | At this time, Windows certificate authentication is not supported. |
| Forms-based authentication | • Lightweight Directory Access Protocol (LDAP)<br>• SQL database or other database | |

| Method | Examples | Notes |
|---|---|---|
| | - Custom or third-party membership and role providers | |
| SAML token-based authentication | - Active Directory Federation Services (AD FS) 2.0<br>- Third-party identity provider<br>- Lightweight Directory Access Protocol (LDAP) | Supported only with SAML 1.1 that uses the WS-Federation Passive profile. |

# Authentication modes — classic or claims-based

SharePoint Server 2010 introduces claims-based authentication, which is built on Windows Identity Foundation (WIF). You can use any of the supported authentication methods with claims-based authentication. Or, you can use classic-mode authentication, which supports Windows authentication.

When you create a Web application, you select one of the two authentication modes to use with the Web application, either claims-based or classic-mode.

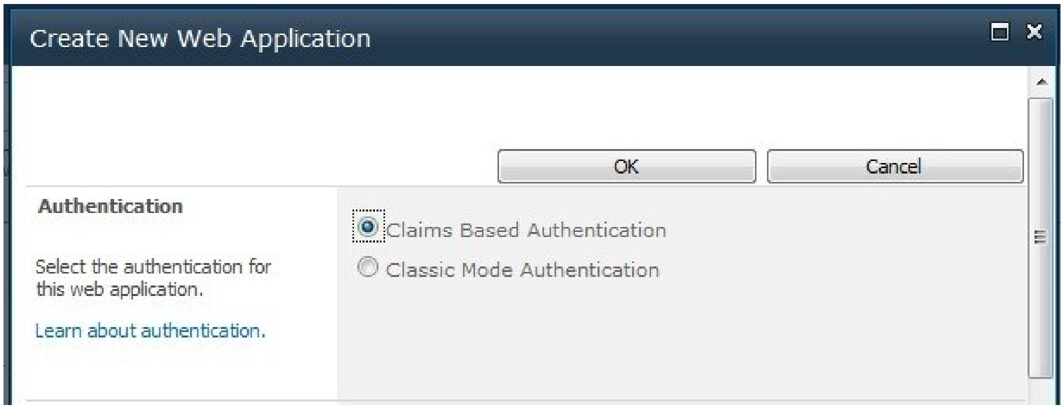

Authentication

Select the authentication for this web application.

Learn about authentication.

● Claims Based Authentication

○ Classic Mode Authentication

If you select classic-mode, you can implement Windows authentication and the user accounts are treated by SharePoint Server 2010 as Active Directory Domain Services (AD DS) accounts.

If you select claims-based authentication, SharePoint Server 2010 automatically changes all user accounts to claims identities, resulting in a claims token for each user. The claims token contains the claims pertaining to the user. Windows accounts are converted into Windows claims. Forms-based membership users are transformed into forms-based authentication claims. Claims that are included in SAML-based tokens can be used by SharePoint Server 2010. Additionally, SharePoint developers and administrators can augment user tokens with additional claims. For example, user Windows accounts and forms-based accounts can be augmented with additional claims that are used by SharePoint Server 2010.

The following chart summarizes the support for authentication types by each authentication mode.

| Type | Classic-mode authentication | Claims-based authentication |
|---|---|---|
| Windows<br>• NTLM<br>• Kerberos<br>• Anonymous<br>• Basic<br>• Digest | Yes | Yes |
| Forms-based authentication<br>• LDAP<br>• SQL database or other database<br>• Custom or third-party membership and role providers | No | Yes |
| SAML token-based authentication<br>• AD FS 2.0<br>• Third-party identity provider<br>• LDAP | No | Yes |

A SharePoint Server 2010 farm can include a mix of Web applications that use both modes. Services do not differentiate between user accounts that are traditional Windows accounts and Windows claims accounts. Consequently, a user who belongs to sites that are configured to use a mix of authentication modes will receive search results that include results from all the sites that the user has access to, regardless of the mode that is configured for Web applications. The user is not interpreted as two different user accounts. This is because services and service applications use claims identities for inter-farm communication regardless of the mode that is selected for Web applications and users.

However, users who belong to more than one user repository that is recognized by SharePoint Server Web applications are treated as separate user accounts, depending on which identity they use to log in.

The following guidance will help you decide which mode to select:

- For new implementations of SharePoint Server 2010, use claims-based authentication. With this option, all supported authentication types are available for Web applications. There is no practical reason to select classic-mode authentication for new deployments, even if your environment includes only Windows accounts. Windows authentication is implemented the same way regardless of the mode that is selected. There are no additional steps to implement Windows authentication when you use the claims-based authentication mode.

- If you are upgrading a previous version solution to SharePoint Server 2010 and the solution includes only Windows accounts, you can use classic-mode authentication. This lets you use the same design for zones and URLs.
- If you are upgrading a solution that requires forms-based authentication, the only option is to upgrade to claims-based authentication.

If you are upgrading from an earlier version to SharePoint Server 2010 and you select claims-based authentication, be aware of the following considerations:

- Custom code might need to be updated. Web Parts or other custom code that relies on or uses Windows identities will have to be updated. If the custom code uses Windows identities, use classic-mode authentication until the code is updated.
- Migrating many Windows users to claims identities takes time. When you change a Web application from classic mode to claims-based during the upgrade process, you must use Windows PowerShell to convert Windows identities to claims identities. This can be a time-consuming process. Be sure to allow enough time during the upgrade process to complete this task.
- Search alerts are currently not supported with claims-based authentication.

Claims authentication is built on WIF. WIF is a set of .NET Framework classes that are used to implement claims-based identity. Claims authentication relies on standards such as WS-Federation, WS-Trust, and protocols such as SAML. For more information about claims authentication, see the following resources:

- Claims-based Identity for Windows: An Introduction to Active Directory Federation Services 2.0, Windows CardSpace 2.0, and Windows Identity Foundation (white paper) (http://go.microsoft.com/fwlink/?LinkId=198942)
- Windows Identity Foundation home page (http://go.microsoft.com/fwlink/?LinkId=198943)

You do not have to be a claims architect to use claims authentication in SharePoint Server 2010. However, implementing SAML token-based authentication requires coordination with administrators of your claims-based environment, as described later in this article.

If you are using Microsoft Unified Access Gateway (UAG) 2010 with Service Pack 1 (SP1) installed, you should be aware of additional capabilities that were added in UAG SP1 that you should be aware of:

- UAG SP1 adds support for Active Directory Federated Services Version 2.0 (ADFS 2.0), and UAG is a claims-aware relying party that now supports publishing with claims-based authentication. (Partner access using single sign-on to applications or to servers running SharePoint Server and that are not claims-aware is still supported.)

  The following steps detail the process flow for authenticating a remote client through a server that is running UAG to a server that is running SharePoint Server:

  a. The remote client connects to a UAG portal instead of directly to SharePoint Server.

b. The server running UAG determines that the client does not have a SAML token.

c. The UAG server creates and then publishes an ADFS application that enables the remote client to access the server running ADFS in order to authenticate.

d. The remote client gets a SAML token from the server running ADFS.

The SAML token contains the credentials of the remote client.

e. The server running UAG uses single sign-on to connect to SharePoint Server.

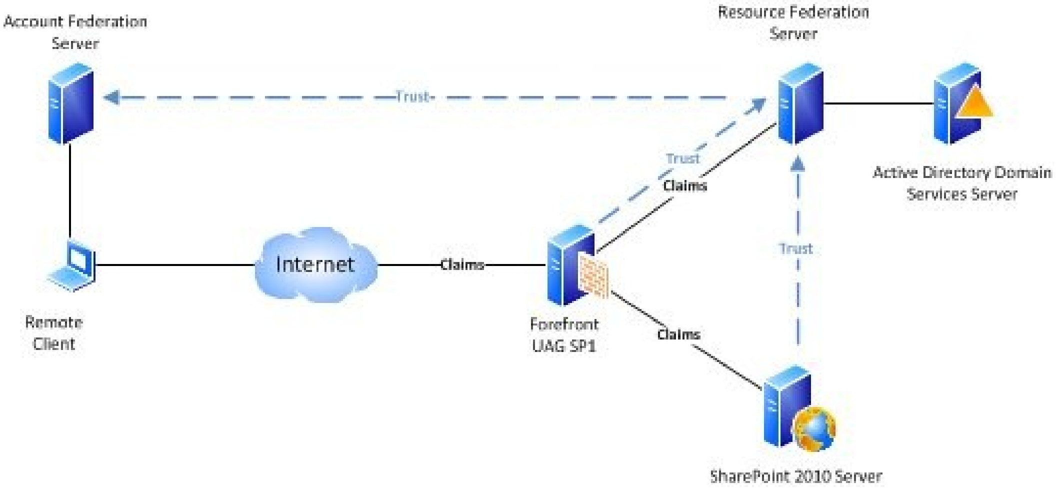

- UAG SP1 adds claims-based authorization. For example: If a user has a claims role, UAG can allow or deny the user's access based on the value of the claim. These rules are set through policy in UAG and are mapped to roles in ADFS.

  **Note:**
  These claims-based authorization rules can only be used when UAG is a relying party of ADFS.

- UAG SP1 adds single sign-out functionality. When users sign out, they are also signed out from all applications that rely on the authenticating federation server.

  There are a few ways that a client can sign out (or be signed out):

  - A user can sign out from the UAG portal.
  - A timed interval of inactivity can sign out a user.
  - Scheduled signoff times of UAG can sign out a user.

  For more information about single sign-out, see Overview of AD FS 2.0 with Forefront UAG (http://go.microsoft.com/fwlink/?LinkId=207207).

- UAG can still provide SSO access if an application uses NTLM or Kerberos authentication, and UAG performs Kerberos translation for clients. For more information, see Configuring single sign-on with Kerberos constrained delegation to non-claims-aware applications (http://go.microsoft.com/fwlink/?LinkId=207208).

# Implementing Windows authentication

The process of implementing Windows authentication methods is similar for both authentication modes (classic or claims-based). Choosing claims-based authentication for a Web application does not increase the complexity of implementing Windows authentication methods. This section summarizes the process for each method.

**Integrated Windows authentication — Kerberos and NTLM**

Both Kerberos protocol and NTLM are Integrated Windows authentication methods, which let clients seamlessly authenticate without being prompted for credentials. Users who access SharePoint sites from Windows Explorer will authenticate by using the credentials the Internet Explorer process is running under. By default, these credentials are the credentials that the user used to log on to the computer. Services or applications that access SharePoint Server in Integrated Windows authentication mode will attempt to authenticate by using the credentials of the running thread, which is the identity of the process by default.

NTLM is the simplest form of Windows authentication to implement. Simply select this option when you are creating a Web application.

Kerberos protocol is a secure protocol that supports ticketing authentication. Use of the Kerberos protocol requires additional configuration of the environment. To enable Kerberos authentication, the client and server computers must have a trusted connection to the domain Key Distribution Center (KDC). Configuring the Kerberos protocol involves setting up service principal names (SPNs) in AD DS before you install SharePoint Server 2010.

The following steps summarize the process of configuring Kerberos authentication:

1. Configure Kerberos authentication for SQL communications by creating SPNs in AD DS for the SQL Server service account.
2. Create SPNs for Web applications that will use Kerberos authentication.
3. Install the SharePoint Server 2010 farm.
4. Configure specific services within the farm to use specific accounts.
5. Create the Web applications that will use Kerberos authentication.

For more information, see Configure Kerberos authentication (SharePoint Server 2010).

Additionally, for claims-authentication Web applications, the claims to Windows token service must be configured for constrained delegation. Constrained delegation is required to convert claims to Windows tokens. For environments that include multiple forests, a two-way trust between forests is required to use the claims to Windows token service. For more information about how to configure this service, see Configure Kerberos authentication for the claims to Windows token service (SharePoint Server 2010).

Kerberos authentication allows delegation of client credentials to access back-end data systems, which requires additional configuration depending on the scenario. The following table provides examples.

| Scenario | Additional configuration |
|---|---|
| Delegating a client's identity to a back-end server.<br>Displaying RSS feeds to authenticated content. | Configure Kerberos constrained delegation for computers and service accounts. |
| Identity delegation for Microsoft SQL Server Reporting Services (SSRS) | Configure SPNs for SQL Server Reporting Services accounts.<br>Configure delegation for SQL Server Reporting Services. |
| Identity delegation for Excel Services in SharePoint | Configure constrained delegation for servers that run Excel Services.<br>Configure constrained delegation for the Excel Services service account. |

For more information about how to configure Kerberos authentication, including configuration steps for common scenarios, see Configuring Kerberos Authentication for Microsoft SharePoint 2010 Products and Technologies (white paper) (http://go.microsoft.com/fwlink/?LinkID=197178).

**Digest and Basic**

Implementing Digest and Basic authentication requires configuring these authentication methods directly in Internet Information Services (IIS).

# Implementing forms-based authentication

Forms-based authentication is an identity management system that is based on ASP.NET membership and role provider authentication. In SharePoint Server 2010, forms-based authentication is available only when you use claims-based authentication.

Forms-based authentication can be used against credentials stored in AD DS, in a database such as a SQL Server database, or in an LDAP data store such as Novell eDirectory, Novell Directory Services (NDS), or Sun ONE. Forms-based authentication enables user authentication based on validation of credential input from a logon form. Unauthenticated requests are redirected to a logon page, where the user must provide valid credentials and submit the form. If the request can be authenticated, the system issues a cookie that contains a key for reestablishing the identity for subsequent requests.

To use forms-based authentication to authenticate users against an identity management system that is not based on Windows or that is external, you must register the membership provider and role manager in the Web.config file. Registering the role manager is a new requirement for SharePoint Server 2010. In the previous version, this was optional. SharePoint Server 2010 uses the standard ASP.NET role manager interface to gather group information about the current user. Each ASP.NET role is treated as a domain group by the authorization process in SharePoint Server 2010. You register role managers in the Web.config file the same way that you register membership providers for authentication.

If you want to manage membership users or roles from the SharePoint Central Administration Web site, you must register the membership provider and the role manager in the Web.config file for the Central Administration Web site. You must also register the membership provider and the role manager in the Web.config file for the Web application that hosts the content.

For more information about how to configure forms-based authentication, see the following resources:

- TechNet article: Configure forms-based authentication for a claims-based Web application (SharePoint Server 2010)
- MSDN blog article: Claims-based authentication "Cheat Sheet" Part 1 (http://go.microsoft.com/fwlink/?LinkId=198944)
- MSDN article: Forms Authentication in SharePoint Products and Technologies (Part 2): Membership and Role Provider Samples (http://go.microsoft.com/fwlink/?LinkId=198945)

# Implementing SAML token-based authentication

SAML token-based authentication requires coordination with administrators of a claims-based environment, whether it is your own internal environment or a partner environment. AD FS 2.0 is an example of a claims-based environment.

A claims-based environment includes an identity provider security token service (IP-STS). The IP-STS issues SAML tokens on behalf of users who are included in the associated user directory. Tokens can include any number of claims about a user, such as a user name and groups the user belongs to.

SharePoint Server 2010 takes advantage of claims that are included in tokens provided by an IP-STS to authorize users. In claims environments, an application that accepts SAML tokens is known as a relying party STS (RP-STS). A relying party application receives the SAML token and uses the claims inside to decide whether to grant the client access to the requested resource. In SharePoint 2010 Products, each Web application that is configured to use a SAML provider is added to the IP-STS server as a separate RP-STS entry. A SharePoint farm can include multiple RP-STS entries.

Implementing SAML token-based authentication with SharePoint 2010 Products involves the following processes that require advance planning:

1. Export the token-signing certificate from the IP-STS. This certificate is known as the ImportTrustCertificate. Copy the certificate to a server computer in the SharePoint Server 2010 farm.
2. Define the claim that will be used as the unique identifier of the user. This is known as the identity claim. Many examples of this process use the user e-mail name as the user identifier. Coordinate with the administrator of the IP-STS to determine the correct identifier because only the owner of the IP-STS knows which value in the token will always be unique per user. Identifying the unique identifier for the user is part of the claims-mapping process. Claims mappings are created by using Windows PowerShell.
3. Define additional claims mappings. Define which additional claims from the incoming token will be used by the SharePoint Server 2010 farm. User roles are an example of a claim that can be used to

permission resources in the SharePoint Server 2010 farm. All claims from an incoming token that do not have a mapping will be discarded.

4. Create a new authentication provider by using Windows PowerShell to import the token-signing certificate. This process creates the SPTrustedIdentityTokenIssuer. During this process, you specify the identity claim and additional claims that you have mapped. You must also create and specify a realm that is associated with the first SharePoint Web applications that you are configuring for SAML token-based authentication. After the SPTrustedIdentityTokenIssuer is created, you can create and add more realms for additional SharePoint Web applications. This is how you configure multiple Web applications to use the same SPTrustedIdentityTokenIssuer.
5. For each realm that is added to the SPTrustedIdentityTokenIssuer, you must create an RP-STS entry on the IP-STS. This can be done before the SharePoint Web application is created. Regardless, you must plan the URL before you create the Web applications.
6. Create a new SharePoint Web application and configure it to use the newly created authentication provider. The authentication provider will appear as an option in Central Administration when claims mode is selected for the Web application.

You can configure multiple SAML token-based authentication providers. However, you can only use a token-signing certificate once in a farm. All providers that are configured will appear as options in Central Administration. Claims from different trusted STS environments will not conflict.

If you are implementing SAML token-based authentication with a partner company and your own environment includes an IP-STS, we recommend that you work with the administrator of your internal claims environment to establish a trust relationship from your internal IP-STS to the partner STS. This approach does not require adding an additional authentication provider to your SharePoint Server 2010 farm. It also allows your claims administrators to manage the whole claims environment.

**Note:**

If you use SAML token-based authentication with AD FS on a SharePoint Server 2010 farm that has multiple Web servers in a load-balanced configuration, there might be an effect on the performance and functionality of client Web-page views. When AD FS provides the authentication token to the client, that token is submitted to SharePoint Server 2010 for each permission-restricted page element. If the load-balanced solution is not using affinity, each secured element is authenticated to more than one SharePoint Server 2010 server, which might result in rejection of the token. After the token is rejected, SharePoint Server 2010 redirects the client to reauthenticate back to the AD FS server. After this occurs, an AD FS server might reject multiple requests that are made in a short time period. This behavior is by design, to protect against a denial of service attack. If performance is adversely affected or pages do not load completely, consider setting network load balancing to single affinity. This isolates the requests for SAML tokens to a single Web server.

For more information about how to configure SAML token-based authentication, see the following resources:

- TechNet article: Configure authentication using a SAML security token (SharePoint Server 2010)

- MSDN blog article: Claims-based authentication "Cheat Sheet" Part 2 (http://go.microsoft.com/fwlink/?LinkId=198946)
- TechNet blog article: Planning Considerations for Claims Based Authentication in SharePoint 2010 (http://go.microsoft.com/fwlink/?LinkId=198947)
- TechNet blog article: Creating both an Identity and Role Claim for a SharePoint 2010 Claims Auth Application (http://go.microsoft.com/fwlink/?LinkId=198948)
- TechNet blog article: How to Create Multiple Claims Auth Web Apps in a Single SharePoint 2010 Farm (http://go.microsoft.com/fwlink/?LinkId=198949)

# Choosing authentication for LDAP environments

LDAP environments can be implemented by using either forms-based authentication or SAML token-based authentication. We recommend that you use forms-based authentication because it is less complex. However, if the environment supports WS-Federation 1.1 and SAML Token 1.1, then SAML is recommended. Profile synchronization is not supported with LDAP providers that are not associated with ADFS 2.0.

# Planning zones for Web applications

Zones represent different logical paths for gaining access to the same sites in a Web application. Each Web application can include as many as five zones. When a Web application is created, the default zone is created. Additional zones are created by extending the Web application and selecting one of the remaining zone names: intranet, extranet, Internet, or custom.

In previous versions, zones are used to implement different types of authentication for users coming from different networks or authentication providers. In the current version, claims authentication allows multiple types of authentication to be implemented on the same zone.

Your plan for zones will depend on which of the following modes is selected for a Web application:

- Classic mode — Similar to previous versions, only one type of authentication can be implemented per zone. However, in the current version,  only Windows authentication can be implemented when classic mode is selected. Consequently, multiple zones can be used only to implement multiple types of Windows authentication, or to implement the same type of Windows authentication against different Active Directory stores.
- Claims authentication —  Multiple authentication providers can be implemented on a single zone. Multiple zones can be used also.

**Implementing more than one type of authentication on a single zone**

If you are using claims authentication and implementing more than one type of authentication, we recommend that you implement multiple types of authentication on the default zone. This results in the same URL for all users.

When you are implementing multiple types of authentication on the same zone, the following restrictions apply:

- Only one instance of forms-based authentication can be implemented on a zone.
- Central Administration allows you to use both an Integrated Windows method and Basic at the same time. Otherwise, more than one type of Windows authentication cannot be implemented on a zone.

If multiple SAML token-based authentication providers are configured for a farm, these will all appear as options when you create a Web application or a new zone. Multiple SAML providers can be configured on the same zone.

The following diagram illustrates multiple types of authentication implemented on the default zone for a partner collaboration site.

Multiple types of authentication implemented on the default zone

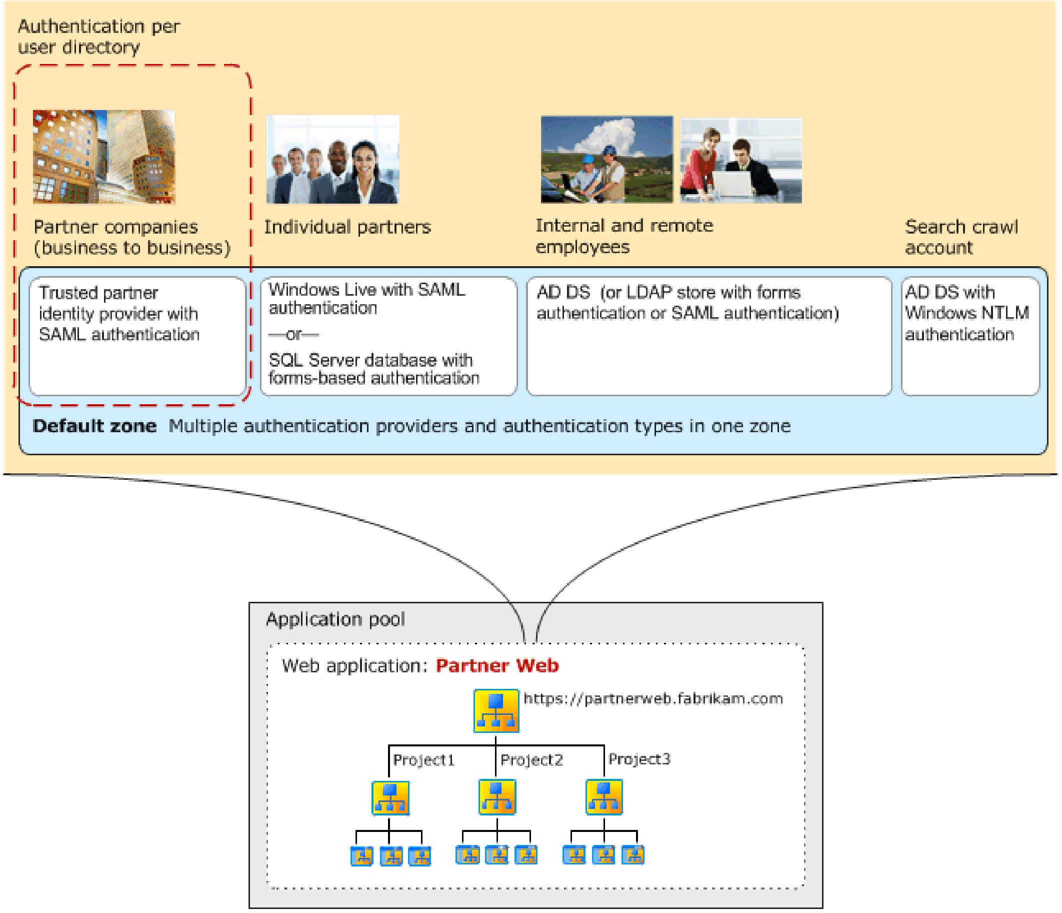
Authentication per user directory
Partner companies (business to business)
Individual partners
Internal and remote employees
Search crawl account
Trusted partner identity provider with SAML authentication
Windows Live with SAML authentication
—or—
SQL Server database with forms-based authentication
AD DS (or LDAP store with forms authentication or SAML authentication)
AD DS with Windows NTLM authentication
Default zone Multiple authentication providers and authentication types in one zone
Application pool
Web application: Partner Web
https://partnerweb.fabrikam.com
Project1
Project2
Project3

In the diagram, users from different directory stores access the partner Web site by using the same URL. A dashed box surrounding partner companies shows the relationship between the user directory and the authentication type that is configured in the default zone. For more information about this design example, see Design sample: Corporate deployment (SharePoint Server 2010).

**Planning for crawling content**

The crawl component requires access to content using NTLM. At least one zone must be configured to use NTLM authentication. If NTLM authentication is not configured on the default zone, the crawl component can use a different zone that is configured to use NTLM authentication.

### Implementing more than one zone

If you plan to implement more than one zone for Web applications, use the following guidelines:

- Use the default zone to implement your most secure authentication settings. If a request cannot be associated with a specific zone, the authentication settings and other security policies of the default zone are applied. The default zone is the zone that is created when you initially create a Web application. Typically, the most secure authentication settings are designed for end-user access. Consequently, end users are likely to access the default zone.
- Use the minimum number of zones that are required to provide access to users. Each zone is associated with a new IIS site and domain for accessing the Web application. Only add new access points when these are required.
- Ensure that at least one zone is configured to use NTLM authentication for the crawl component. Do not create a dedicated zone for the index component unless it is necessary.

The following diagram illustrates multiple zones that are implemented to accommodate different authentication types for a partner collaboration site.

One zone per authentication type

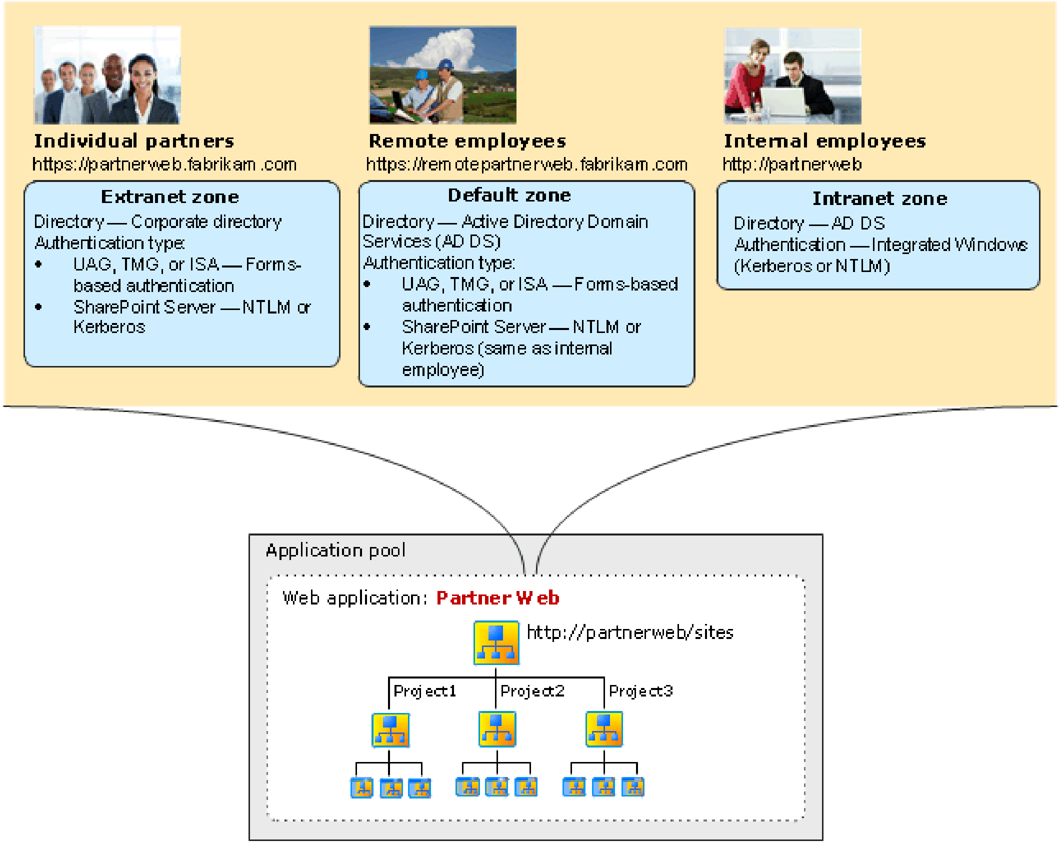

**Individual partners**
https://partnerweb.fabrikam.com

**Extranet zone**
Directory — Corporate directory
Authentication type:
- UAG, TMG, or ISA — Forms-based authentication
- SharePoint Server — NTLM or Kerberos

**Remote employees**
https://remotepartnerweb.fabrikam.com

**Default zone**
Directory — Active Directory Domain Services (AD DS)
Authentication type:
- UAG, TMG, or ISA — Forms-based authentication
- SharePoint Server — NTLM or Kerberos (same as internal employee)

**Internal employees**
http://partnerweb

**Intranet zone**
Directory — AD DS
Authentication — Integrated Windows (Kerberos or NTLM)

In the diagram, the default zone is used for remote employees. Each zone has a different URL associated with it. Employees use a different zone depending on whether they are working in the office or are working remotely. For more information about this design example, see Design sample: Corporate deployment (SharePoint Server 2010).

# Architecture for SAML token-based providers

The architecture for implementing SAML token-based providers includes the following components:

**SharePoint security token service** This service creates the SAML tokens that are used by the farm. The service is automatically created and started on all servers in a server farm. The service is used for inter-farm communication because all inter-farm communication uses claims authentication. This service is also used for authentication methods that are implemented for Web applications that use claims authentication, including Windows authentication, forms-based authentication, and SAML token-based authentication. You must configure the security token service during the deployment process. For more information, see Configure the security token service (SharePoint Server 2010).

**Token-signing certificate (ImportTrustCertificate)** This is the certificate that is exported from an IP-STS. The certificate is copied to one server in the farm. Once you use this certificate to create an SPTrustedIdentityTokenIssuer, you cannot use it again to create another one. If you want to use the certificate to create a different SPTrustedIdentityTokenIssuer, you must delete the existing one first. Before you delete an existing one, you must disassociate it from any Web applications that may be using it.

**Identity claim** The identity claim is the claim from a SAML token that is the unique identifier of the user. Only the owner of the IP-STS knows which value in the token will always be unique for each user. The identity claim is created as a regular claims mapping during the process of mapping all desired claims. The claim that serves as the identity claim is declared when the SPTrustedIdentityTokenIssuer is created.

**Other claims** These claims consist of additional claims from a SAML ticket that describe users. These can include user roles, user groups, or other kinds of claims such as age. All claims mappings are created as objects that are replicated across the servers in a SharePoint Server 2010 farm.

**Realm** In the SharePoint claims architecture, the URI or URL that is associated with a SharePoint Web application that is configured to use a SAML token-based provider represents a realm. When you create a SAML-based authentication provider on the farm, you specify the realms, or Web application URLs, that you want the IP-STS to recognize, one at a time. The first realm is specified when you create the SPTrustedIdentityTokenIssuer. Additional realms can be added after the SPTrustedIdentityTokenIssuer is created. Realms are specified by using syntax similar to the following: $realm = "urn:sharepoint:mysites". After you add the realm to the SPTrustedIdentityTokenIssuer, you must create an RP-STS trust with the realm on the IP-STS server. This process involves specifying the URL for the Web application.

**SPTrustedIdentityTokenIssuer** This is the object that is created on the SharePoint farm that includes the values necessary to communicate with and receive tokens from the IP-STS. When you create the

SPTrustedIdentityTokenIssuer, you specify which token-signing certificate to use, the first realm, the claim that represents the identity claim, and any additional claims. You can only associate a token-signing certificate from an STS with one SPTrustedIdentityTokenIssuer. However, after you create the SPTrustedIdentityTokenIssuer, you can add more realms for additional Web applications. After a realm is added to the SPTrustedIdentityTokenIssuer, it must also be added to the IP-STS as a relying party. The SPTrustedIdentityTokenIssuer object is replicated across servers in the SharePoint Server 2010 farm.

**Relying party security token service (RP-STS)** In SharePoint Server 2010, each Web application that is configured to use a SAML provider is added to the IP-STS server as an RP-STS entry. A SharePoint Server 2010 farm can include multiple RP-STS entries.

**Identity provider security token service (IP-STS)** This is the secure token service in the claims environment that issues SAML tokens on behalf of users who are included in the associated user directory.

The following diagram illustrates the SharePoint 2010 Products claims architecture.

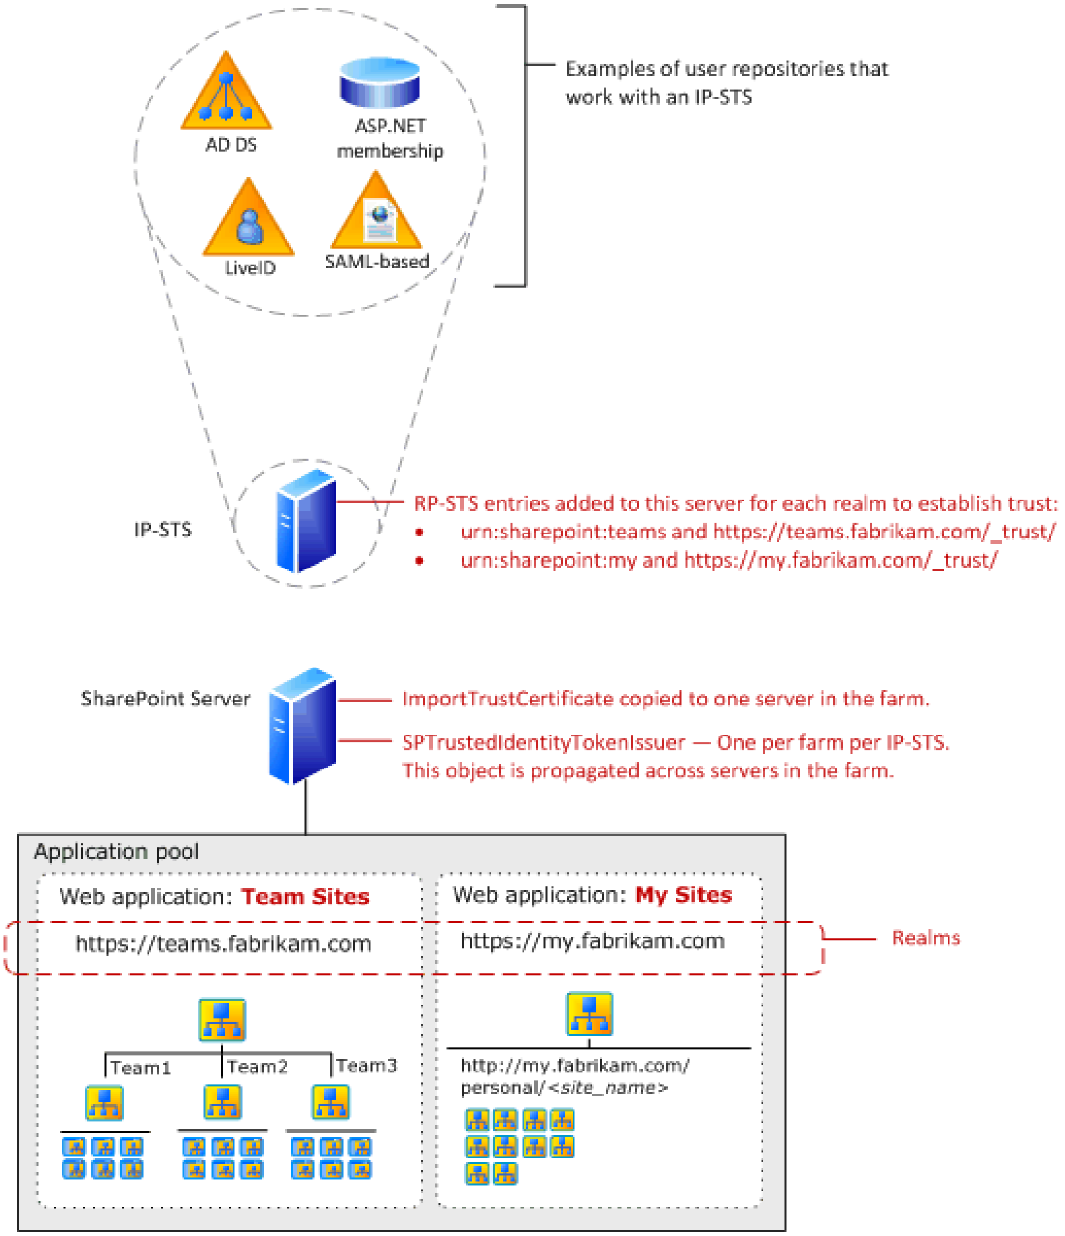
AD DS
ASP.NET membership
LiveID
SAML-based
Examples of user repositories that work with an IP-STS
IP-STS
RP-STS entries added to this server for each realm to establish trust:
• urn:sharepoint:teams and https://teams.fabrikam.com/_trust/
• urn:sharepoint:my and https://my.fabrikam.com/_trust/
SharePoint Server
ImportTrustCertificate copied to one server in the farm.
SPTrustedIdentityTokenIssuer — One per farm per IP-STS. This object is propagated across servers in the farm.
Application pool
Web application: Team Sites
Web application: My Sites
https://teams.fabrikam.com
https://my.fabrikam.com
Realms
Team1
Team2
Team3
http://my.fabrikam.com/personal/<site_name>

The SPTrustedIdentityTokenIssuer object is created by using several parameters. The following diagram illustrates the key parameters.

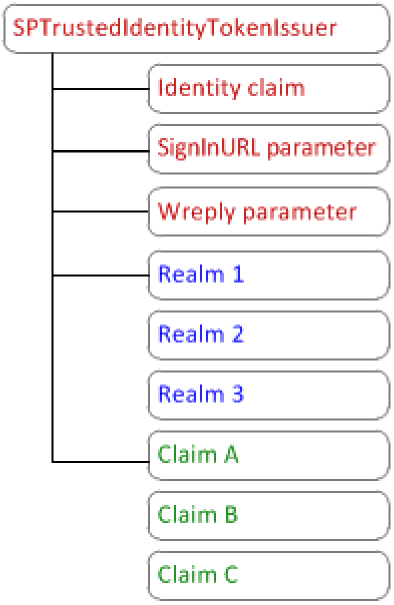

As the diagram illustrates, an SPTrustedIdentityTokenIssuer can include only one identity claim, one SignInURL parameter, and one Wreply parameter. However, it can include multiple realms and multiple claims mappings. The SignInURL parameter specifies the URL to redirect a user request to in order to authenticate to the IP-STS. Some IP-STS servers require the Wreply parameter, which is set to either true or false and is false by default. Only use the Wreply parameter if it is required by the IP-STS.

# Plan the Secure Store Service (SharePoint Server 2010)

In Microsoft SharePoint Server 2010, the Secure Store Service replaces the single sign-on (SSO) feature. The Secure Store Service is a claims-aware authorization service that includes a secure database for storing credentials that are associated with application IDs. These application IDs can be used to authorize access to external data sources.

In this article:

- About the Secure Store Service
- Secure store service preparation
- Application IDs
- Secure store service mappings
- Secure store service and claims authentication

## About the Secure Store Service

The Secure Store Service is an authorization service that runs on an application server. The Secure Store Service provides a database that is used to store credentials (consisting of a user identity and password) for application IDs that can be used by applications to authorize access to shared resources. For example, SharePoint Server 2010 can use the secure store database to store and retrieve credentials for access to external data sources. The Secure Store Service provides support for storing the credentials of multiple back-end systems using multiple application IDs.

## Secure store service preparation

When you prepare to deploy the Secure Store Service, be aware of the following important guidelines:

- Run the Secure Store Service in a separate application pool that is not used for any other service.
- Run the Secure Store Service on a separate application server that is not used for any other service.
- Create the secure store database on a separate application server running SQL Server. Do not use the same SQL Server installation that contains content databases.
- Before you generate a new encryption key, back up the secure store database. You should also back up the secure store database after it is initially created, and again each time credentials are reencrypted. When a new key is generated, the credentials can be re-encrypted with the new key. If the key refresh fails, or the passphrase is forgotten, the credentials will not be useable.
- Back up the encryption key after initially setting up the Secure Store Service, and back up the key again each time it is regenerated.

- Do not store the backup media for the encryption key in the same location as the backup media for the secure store database. If a user obtains a copy of both the database and the key, the credentials stored in the database could be compromised.

# Application IDs

Each Secure Store Service entry contains an application ID that is used to retrieve a set of credentials from the secure store database. Each application ID can have permissions applied so that only specific users or groups can access the credentials that are stored for the application ID. Applications use application IDs to retrieve credentials from the secure store database on behalf of a user. The application can then use the retrieved credentials to access a data source.

Application IDs are used to map users to credential sets. Mappings are available for groups or individuals. In a group mapping, every user who is a member of a specific domain group is mapped to the same set of credentials. In an individual mapping, each individual user is mapped to a unique set of credentials.

# Secure store service mappings

The Secure Store Service supports individual mappings and group mappings. The Secure Store Service maintains a set of credentials for the application IDs of resources that are stored in the secure store database. Individual credentials for an application are retrieved based on the application ID. Individual mappings are useful if you need logging information about individual user access to shared resources. For group mappings, a security layer checks group credentials for multiple domain users against a single set of credentials for a resource identified by an application ID that is stored in the secure store database. Group mappings are easier to maintain than individual mappings, and can provide improved performance.

# Secure store service and claims authentication

The Secure Store Service is a claims-aware service. It can accept security tokens and decrypt them to get the application ID, and then perform a lookup.. When a SharePoint Server 2010 Security Token Service (STS) issues a security token in response to an authentication request, the Secure Store Service decrypts the token and reads the application ID value. The Secure Store Service uses the application ID to retrieve credentials from the secure store database. The credentials are then used to authorize access to resources.

**See Also**

Configure the Secure Store Service (SharePoint Server 2010)

# Plan security hardening (SharePoint Server 2010)

This article describes security hardening for Microsoft SharePoint Server 2010 Web server, application server, and database server roles, and gives detailed guidance about the specific hardening requirements for ports, protocols, and services in Microsoft SharePoint 2010 Products.

In this article:

- Secure server snapshots
- Specific port, protocol, and service guidance

## Secure server snapshots

In a server farm environment, individual servers play specific roles. Security hardening recommendations for these servers depend on the role each server plays. This article contains secure snapshots for two categories of server roles:

- Web server and application server roles
- Database server role

The snapshots are divided into common configuration categories. The characteristics defined for each category represent the optimal hardened state for Microsoft SharePoint 2010 Products. This article does not include hardening guidance for other software in the environment.

### Web server and application server roles

This section identifies hardening characteristics for Web servers and application servers. Some of the guidance applies to specific service applications; in these cases, the corresponding characteristics need to be applied only on the servers that are running the services associated with the specified service applications.

| Category | Characteristic |
|---|---|
| Services listed in the Services MMC snap-in | Enable the following services:<br>• File and Printer Sharing<br>• ASP.NET State service (if you are using InfoPath Forms Services or Microsoft Project Server 2010)<br>• View State service (if you are using InfoPath Forms Services) |

| | |
|---|---|
| | • World Wide Web Publishing Service<br>Ensure that these services are not disabled:<br>• Claims to Windows Token Service<br>• SharePoint 2010 Administration<br>• SharePoint 2010 Timer<br>• SharePoint 2010 Tracing<br>• SharePoint 2010 VSS Writer<br>Ensure that these services are not disabled on the servers that host the corresponding roles:<br>• SharePoint 2010 User Code Host<br>• SharePoint Foundation Search V4<br>• SharePoint Server Search 14<br>• The following services are required by the User Profile service application on the server that imports profiles from the directory store:<br>    • Forefront Identity Manager service<br>    • Forefront Identity Manager Synchronization service |
| Ports and protocols | • TCP 80, TCP 443 (SSL)<br>• Custom ports for search crawling, if configured<br>• File and Printer Sharing service —either of the following, used by search roles:<br>    • Direct-hosted SMB (TCP/UDP 445) — this is the recommended port<br>    • NetBIOS over TCP/IP (NetBT) (TCP/UDP ports 137, 138, 139) — disable this port if you do not  use it<br>• Ports required for communication between Web servers and service applications (the default is HTTP):<br>    • HTTP binding: 32843<br>    • HTTPS binding: 32844<br>    • net.tcp binding: 32845 (only if a third party has implemented this option for a service application)<br>• Ports required for synchronizing profiles between SharePoint 2010 Products and Active |

| | |
|---|---|
| | Directory on the server that runs the Forefront Identity Management agent:<br>• TCP/5725<br>• TCP/UDP 389 (LDAP service)<br>• TCP/UDP 88 (Kerberos)<br>• TCP/UDP 53 (DNS)<br>• UDP 464 (Kerberos Change Password)<br>For information about how to synchronize profiles with other directory stores, see User Profile service hardening requirements, later in this article.<br>• UDP port 1434 and TCP port 1433 — default ports for SQL Server communication. If these ports are blocked on the SQL Server computer (recommended) and databases are installed on a named instance, configure a SQL Server client alias for connecting to the named instance.<br>• TCP/IP 32846 for the Microsoft SharePoint Foundation User Code Service (for sandbox solutions) — This port must be open for outbound connections on all Web servers. This port must be open for inbound connections on Web servers or application servers where this service is turned on.<br>• Ensure that ports remain open for Web applications that are accessible to users.<br>• Block external access to the port that is used for the Central Administration site.<br>• TCP/25 (SMTP for e-mail integration) |
| Registry | No additional guidance |
| Auditing and logging | If log files are relocated, ensure that the log file locations are updated to match. Update directory access control lists (ACLs) also. |
| Code access security | Ensure that you have a minimal set of code access security permissions enabled for your Web application. The <trust> element in the Web.config file for each Web application should be set to |

| | |
|---|---|
| | WSS_Minimal (where WSS_Minimal has its low defaults as defined in 14\config\wss_minimaltrust.config or by your own custom policy file, which is minimally set.) |
| Web.config | Follow these recommendations for each Web.config file that is created after you run Setup:<br>• Do not allow compilation or scripting of database pages via the PageParserPaths elements.<br>• Ensure <SafeMode> CallStack=""false"" and AllowPageLevelTrace=""false"".<br>• Ensure that the Web Part limits around maximum controls per zone is set low.<br>• Ensure that the SafeControls list is set to the minimum set of controls needed for your sites.<br>• Ensure that your Workflow SafeTypes list is set to the minimum level of SafeTypes needed.<br>• Ensure that customErrors is turned on (<customErrors mode=""On""/>).<br>• Consider your Web proxy settings as needed (<system.net>/<defaultProxy>).<br>• Set the Upload.aspx limit to the highest size you reasonably expect users to upload (default is 2 GB). Performance can be affected by uploads that exceed 100 MB. |

## Database server role

The primary recommendation forSharePoint 2010 Products is to secure inter-farm communication by blocking the default ports used for Microsoft SQL Server communication and establishing custom ports for this communication instead. For more information about how to configure ports for SQL Server communication, see Blocking the standard SQL Server ports, later in this article.

| Category | Characteristic |
|---|---|
| Ports | • Block UDP port 1434.<br>• Consider blocking TCP port 1433. |

This article does not describe how to secure SQL Server. For more information about how to secure SQL Server, see Securing SQL Server (http://go.microsoft.com/fwlink/?LinkId=186828).

# Specific port, protocol, and service guidance

The rest of this article describes in greater detail the specific hardening requirements for SharePoint 2010 Products.

In this section:

- Blocking the standard SQL Server ports
- Service application communication
- File and Printer Sharing service requirements
- User Profile service hardening requirements
- Connections to external servers
- Service requirements for e-mail integration
- Service requirements for session state
- SharePoint 2010 Products services
- Web.config file

## Blocking the standard SQL Server ports

The specific ports used to connect to SQL Server are affected by whether databases are installed on a default instance of SQL Server or a named instance of SQL Server. The default instance of SQL Server listens for client requests on TCP port 1433. A named instance of SQL Server listens on a randomly assigned port number. Additionally, the port number for a named instance can be reassigned if the instance is restarted (depending on whether the previously assigned port number is available).

By default, client computers that connect to SQL Server first connect by using TCP port 1433. If this communication is unsuccessful, the client computers query the SQL Server Resolution Service that is listening on UDP port 1434 to determine the port on which the database instance is listening.

The default port-communication behavior of SQL Server introduces several issues that affect server hardening. First, the ports used by SQL Server are well-publicized ports and the SQL Server Resolution Service has been the target of buffer overrun attacks and denial-of-service attacks, including the "Slammer" worm virus. Even if SQL Server is updated to mitigate security issues in the SQL Server Resolution Service, the well-publicized ports remain a target. Second, if databases are installed on a named instance of SQL Server, the corresponding communication port is randomly assigned and can change. This behavior can potentially prevent server-to-server communication in a hardened environment. The ability to control which TCP ports are open or blocked is essential to securing your environment.

Consequently, the recommendation for a server farm is to assign static port numbers to named instances of SQL Server and to block UDP port 1434 to prevent potential attackers from accessing the

SQL Server Resolution Service. Additionally, consider reassigning the port used by the default instance and blocking TCP port 1433.

There are several methods you can use to block ports. You can block these ports by using a firewall. However, unless you can be sure that there are no other routes into the network segment and that there are no malicious users that have access to the network segment, the recommendation is to block these ports directly on the server that hosts SQL Server. This can be accomplished by using Windows Firewall in Control Panel.

### Configuring SQL Server database instances to listen on a nonstandard port

SQL Server provides the ability to reassign the ports that are used by the default instance and any named instances. In SQL Server 2005 and SQL Server 2008, you reassign ports by using SQL Server Configuration Manager.

### Configuring SQL Server client aliases

In a server farm, all front-end Web servers and application servers are SQL Server client computers. If you block UDP port 1434 on the SQL Server computer, or you change the default port for the default instance, you must configure a SQL Server client alias on all servers that connect to the SQL Server computer.

To connect to an instance of SQL Server 2005 or SQL Server 2008, you install SQL Server client components on the target computer and then configure the SQL Server client alias by using SQL Server Configuration Manager. To install SQL Server client components, run Setup and select only the following client components to install:

- Connectivity Components
- Management Tools (includes SQL Server Configuration Manager)

For specific hardening steps for blocking the standard SQL ports, see Harden SQL Server for SharePoint environments (SharePoint Server 2010).

## Service application communication

By default, communication between Web servers and service applications within a farm takes place by using HTTP with a binding to port 32843. When you publish a service application, you can select either HTTP or HTTPS with the following bindings:

- HTTP binding: port 32843
- HTTPS binding: port 32844

Additionally, third parties that develop service applications can implement a third choice:

- net.tcp binding: port 32845

You can change the protocol and port binding for each service application. On the Service Applications page in Central Administration, select the service application, and then click **Publish**.

Communication between service applications and SQL Server takes place over the standard SQL Server ports or the ports that you configure for SQL Server communication.

## File and Printer Sharing service requirements

Several core features depend on the File and Printer Sharing service and the corresponding protocols and ports. These include, but are not limited to, the following:

- **Search queries** All search queries require the File and Printer Sharing service.
- **Crawling and indexing content** To crawl content, servers that include crawl components send requests through the front-end Web server. The front-end Web server communicates with content databases directly and sends results back to the servers that include crawl components. This communication requires the File and Printer Sharing service.
- **Index propagation** If a Search service application is configured with crawl components and query components that are distributed across multiple servers, the servers with crawl components copy content index files to the servers with query components. This action requires the File and Printer Sharing service and its corresponding protocols and ports.

The File and Printer Sharing service requires the use of named pipes. Named pipes can communicate by using either direct-hosted SMB or NetBT protocols. For a secure environment, direct-hosted SMB is recommended instead of NetBT. The hardening recommendations provided in this article assume that SMB is used.

The following table describes the hardening requirements that are introduced by the dependency on the File and Printer Sharing service.

| Category | Requirements | Notes |
|---|---|---|
| Services | File and Printer Sharing | Requires the use of named pipes. |
| Protocols | Named pipes that use direct-hosted SMB<br>Disable NetBT | Named pipes can use NetBT instead of direct-hosted SMB. However, NetBT is not considered as secure as direct-hosted SMB. |
| Ports | Either of the following:<br>• Direct-hosted SMB (TCP/UDP 445) — recommended<br>• NetBT (TCP/UDP ports 137, 138, 139) | Disable NetBT (ports 137, 138, and 139) if it is not being used |

For more information about how to disable NetBT, see the Microsoft Knowledge Base article 204279, Direct hosting of SMB over TCP/IP (http://go.microsoft.com/fwlink/?LinkId=76143).

## User Profile service hardening requirements

The User Profile service application uses the Forefront Identity Management agent to synchronize profiles between SharePoint 2010 Products and Active Directory or a Lightweight Directory Access Protocol (LDAP) directory service. The Forefront Identity Management agent is installed on all servers in a SharePoint farm, but is only required on the server that is set up to synchronize with the directory store.

The Forefront Identity Management agent includes the following two services that must remain enabled on the server that is set up to crawl Active Directory or another directory store:

- Forefront Identity Manager service
- Forefront Identity Manager Synchronization service

Additionally, TCP port 5725 must be open on the server that runs the Forefront Identity Management agent and is set up to crawl a directory store.

In Active Directory environments, the following ports must remain open for communication between the SharePoint 2010 Products server that synchronizes with the directory store and the server that is running Active Directory:

- TCP/UDP 389 (LDAP service)
- TCP/UDP 88 (Kerberos)
- TCP/UDP 53 (DNS)
- UDP 464 (Kerberos Change Password)

For more information about hardening requirements for the Forefront Identity Management agent, including port requirements for other directory types, see Management Agent Communication Ports, Rights, and Permissions (http://go.microsoft.com/fwlink/?LinkId=186832).

## Connections to external servers

Several features of SharePoint Server 2010 can be configured to access data that resides on server computers outside of the server farm. If you configure access to data that is located on external server computers, ensure that you enable communication between the appropriate computers. In most cases, the ports, protocols, and services that are used depend on the external resource. For example:

- Connections to file shares use the File and Printer Sharing service.
- Connections to external SQL Server databases use the default or customized ports for SQL Server communication.
- Connections to Oracle databases typically use OLE DB.
- Connections to Web services use both HTTP and HTTPS.

The following table lists features that can be configured to access data that resides on server computers outside the server farm.

| Feature | Description |
|---|---|
| Content crawling | You can configure crawl rules to crawl data that resides on external resources, including Web sites, file shares, Exchange public folders, and business data applications. When crawling external data sources, the crawl role communicates directly with these external resources.<br>For more information, see Plan to crawl content (Office SharePoint Server) [ http://technet.microsoft.com/en-us/library/cc262926.aspx ] . |
| Business Data Connectivity connections | Web servers and application servers communicate directly with computers that are configured for Business Data Connectivity connections.<br>For more information, see Plan for business data connections with the Business Data Catalog [ http://technet.microsoft.com/en-us/library/cc263252.aspx ] . |
| Receiving Microsoft Office Excel workbooks | If workbooks opened in Excel Services Application connect to any external data sources (for example, Analysis Services and SQL Server), appropriate TCP/IP ports need to be opened for connecting to these external data sources. For more information, see Plan external data connections for Excel Services [ http://technet.microsoft.com/en-us/library/cc262899.aspx ] .<br>If Universal Naming Convention (UNC) paths are configured as trusted locations in Excel Services Application, the Excel Calculation Services application role uses the protocols and ports used by the File and Printer Sharing service to receive Office Excel workbooks over a UNC path.<br>Workbooks that are stored in content databases or that are uploaded or downloaded from sites by users are not affected by this communication. |

## Service requirements for e-mail integration

E-mail integration requires the use of two services:

- SMTP service
- Microsoft SharePoint Directory Management service

### SMTP service

E-mail integration requires the use of the Simple Mail Transfer Protocol (SMTP) service on at least one of the front-end Web servers in the server farm. The SMTP service is required for incoming e-mail. For outgoing e-mail, you can either use the SMTP service or route outgoing email through a dedicated e-mail server in your organization, such as a Microsoft Exchange Server computer.

### Microsoft SharePoint Directory Management service

SharePoint 2010 Products include an internal service, the Microsoft SharePoint Directory Management Service, for creating e-mail distribution groups. When you configure e-mail integration, you have the option to enable the Directory Management Service feature, which lets users create distribution lists. When users create a SharePoint group and they select the option to create a distribution list, the Microsoft SharePoint Directory Management Service creates the corresponding Active Directory distribution list in the Active Directory environment.

In security-hardened environments, the recommendation is to restrict access to the Microsoft SharePoint Directory Management Service by securing the file associated with this service, which is SharePointEmailws.asmx. For example, you might allow access to this file by the server farm account only.

Additionally, this service requires permissions in the Active Directory environment to create Active Directory distribution list objects. The recommendation is to set up a separate organizational unit (OU) in Active Directory for SharePoint 2010 Products objects. Only this OU should allow write access to the account that is used by the Microsoft SharePoint Directory Management Service.

## Service requirements for session state

Both Project Server 2010 and InfoPath Forms Services maintain session state. If you are deploying these features or products within your server farm, do not disable the ASP.NET State service. Additionally, if you are deploying InfoPath Forms Services, do not disable the View State service.

## SharePoint 2010 Products services

Do not disable services that are installed by SharePoint 2010 Products (listed in the snapshot previously).

If your environment disallows services that run as a local system, you can consider disabling the SharePoint 2010 Administration service only if you are aware of the consequences and can work around them. This service is a Win32 service that runs as a local system.

This service is used by the SharePoint 2010 Timer service to perform actions that require administrative permissions on the server, such as creating Internet Information Services (IIS) Web sites, deploying code, and stopping and starting services. If you disable this service, you cannot complete deployment-related tasks from the Central Administration site. You must use Windows PowerShell to run the Start-SPAdminJob cmdlet (or use the Stsadm.exe command-line tool to run the **execadmsvcjobs** operation) to complete multiple-server deployments for SharePoint 2010 Products and to run other deployment-related tasks.

## Web.config file

The .NET Framework, and ASP.NET in particular, use XML-formatted configuration files to configure applications. The .NET Framework relies on configuration files to define configuration options. The configuration files are text-based XML files. Multiple configuration files can, and typically do, exist on a single system.

System-wide configuration settings for the .NET Framework are defined in the Machine.config file. The Machine.config file is located in the %SystemRoot%\Microsoft.NET\Framework\%VersionNumber%\CONFIG\ folder. The default settings that are contained in the Machine.config file can be modified to affect the behavior of applications that use the .NET Framework on the whole system.

You can change the ASP.NET configuration settings for a single application if you create a Web.config file in the root folder of the application. When you do this, the settings in the Web.config file override the settings in the Machine.config file.

When you extend a Web application by using Central Administration, SharePoint 2010 Products automatically create a Web.config file for the Web application.

The Web server and application server snapshot presented earlier in this article lists recommendations for configuring Web.config files. These recommendations are intended to be applied to each Web.config file that is created, including the Web.config file for the Central Administration site.

For more information about ASP.NET configuration files and editing a Web.config file, see ASP.NET Configuration (http://go.microsoft.com/fwlink/?LinkID=73257).

# Plan automatic password change (SharePoint Server 2010)

To simplify password management, the automatic password change feature enables you to update and deploy passwords without having to perform manual password update tasks across multiple accounts, services, and Web applications. You can configure the automatic password change feature to determine if a password is about to expire and reset the password using a long, cryptographically-strong random string. To implement the automatic password change feature, you have to configure managed accounts.

In this article:

- Configuring managed accounts
- Resetting passwords automatically on a schedule
- Detecting password expiration
- Resetting the account password immediately
- Synchronizing SharePoint Foundation account passwords with Active Directory Domain Services
- Resetting all passwords immediately
- Credential change process

## Configuring managed accounts

Microsoft SharePoint Server 2010 supports the creation of managed accounts to improve security and ensure application isolation. Using managed accounts, you can configure the automatic password change feature to deploy passwords across all services in the farm. You can configure SharePoint Web applications and services, running on application servers in a SharePoint farm, to use different domain accounts. You can create multiple accounts in Active Directory Domain Services (AD DS), and then register each of these accounts in SharePoint Server 2010. You can map managed accounts to various services and Web applications in the farm.

## Resetting passwords automatically on a schedule

Prior to the implementation of the automatic password change feature, updating passwords required resetting each account password in AD DS and then manually updating account passwords on all of the services running on all the computers in the farm. To do this, you had to run the Stsadm command-line tool or use the SharePoint Central Administration Web application. Using the automatic password change feature, you can now register managed accounts and enable SharePoint Server 2010 to control account passwords. Users have to be notified about planned password changes and related service interruptions, but the accounts used by a SharePoint farm, Web applications, and various services can be automatically reset and deployed within the farm as necessary, based on individually configured password reset schedules.

## Detecting password expiration

IT departments typically impose a policy requiring that all domain account passwords be reset on a regular basis, for example, every 60 days. SharePoint Server 2010 can be configured to detect imminent password expiration, and send an e-mail notification to a designated administrator. Even without administrator intervention, SharePoint Server 2010 can be configured to generate and reset passwords automatically. The automatic password reset schedule is also configurable to ensure that the impact of possible service interruptions during a password reset will be minimal.

## Resetting the account password immediately

You can always override any automatic password reset schedule and force an immediate service account password reset, using a specific password value. In this scenario, the password for the service account can also be changed in AD DS by SharePoint Server 2010. The new password is then immediately propagated to other servers in the farm.

## Synchronizing SharePoint Foundation account passwords with Active Directory Domain Services

If AD DS and SharePoint Server 2010 account passwords are not synchronized, services in the SharePoint farm will not start. If an Active Directory administrator changes an Active Directory account password without coordinating the password change with a SharePoint administrator, there is a risk of service interruptions. In this scenario, a SharePoint administrator can immediately reset the password from the Account Management page using the password value that was changed in AD DS. The password is updated and immediately propagated to the other servers in the SharePoint farm.

## Resetting all passwords immediately

If an administrator suddenly leaves your organization, or if the service account passwords need to be immediately reset for any other reason, you can quickly create a Windows PowerShell script that calls the password change cmdlets. You can use the script to generate new random passwords and deploy the new passwords immediately.

## Credential change process

When SharePoint Server 2010 changes the credentials for a managed account, the credential change process will occur on one server in the farm. Each server in the farm will be notified that the credentials are about to change and servers can perform critical pre-change actions, if necessary. If the account password has not yet been changed, then SharePoint Server 2010 will attempt to change the password using either a manually entered password, or a long, cryptographically-strong random string. The complexity settings will be queried from the appropriate policy (network or local), and the generated password will be equivalent to the detected settings.SharePoint Server 2010 will attempt to commit a password change. If it is unable to commit the password change, it will retry, using a new sequence, for a specified number of times. If the account password update process succeeds, it will proceed to the next dependent service, where it will again attempt to commit a password change. If it does not ultimately succeed, each dependent service will be notified that they can resume normal activity. Either success in committing a password change or failure to commit will result in the generation of an automated password change status notification that will be sent by e-mail to farm administrators.

**See Also**

Configure automatic password change (SharePoint Server 2010)

# SQL Server and storage (SharePoint Server 2010)

This section describes how to plan for Microsoft SQL Server and storage configuration for Microsoft SharePoint Server 2010. In this section:

- Overview of SQL Server in a SharePoint environment (SharePoint Server 2010)

  This article describes the relationship between SharePoint Server 2010 and supported versions of SQL Server. It also describes how you can interact with the databases, and introduces ways of using the reporting and business intelligence (BI) features of SQL Server with SharePoint Server 2010.

- SQL Server 2008 R2 and SharePoint 2010 Products: Better Together (white paper) (SharePoint Server 2010)

  This article describes the benefits of using SharePoint Server 2010 with SQL Server 2008 R2 Enterprise Edition.

- Storage and SQL Server capacity planning and configuration (SharePoint Server 2010)

  This article describes a process for planning storage and SQL Server capacity for a SharePoint Server 2010 deployment.

- Overview of Remote BLOB Storage (SharePoint Server 2010)

  This article describes how SharePoint Server 2010 works with remote BLOB storage.

- Plan for Remote BLOB Storage (RBS) (SharePoint Server 2010)

  This article describes the factors to consider when moving to a remote BLOB storage solution.

# Overview of SQL Server in a SharePoint environment (SharePoint Server 2010)

This article describes the relationship between Microsoft SharePoint Server 2010 and supported versions of Microsoft SQL Server. It also describes how you can interact with the databases, and it introduces ways of using the reporting and business intelligence (BI) features of SQL Server with SharePoint Server 2010.

For more information about the supported versions of SQL Server, see Hardware and software requirements (SharePoint Server 2010).

In this article:

- SharePoint 2010 Products and the SQL Server database engine
- SQL Server as a data platform for business intelligence in SharePoint 2010 Products
- SharePoint Server 2010 authoring and publishing tools for business intelligence

## SharePoint 2010 Products and the SQL Server database engine

SharePoint Server 2010 is an application that is built on the SQL Server database engine. Most content and settings in SharePoint Server 2010 are stored in relational databases. SharePoint Server 2010 uses the following kinds of databases:

- **Configuration**   The Configuration database and Central Administration content database are called *configuration databases*. They contain data about farm settings such as the databases used, Internet Information Services (IIS) Web sites or web applications, solutions, Web Part packages, site templates, default quota, and blocked file types. A farm can only have one set of configuration databases.
- **Content**   Content databases store all site content:  site documents, such as files in document libraries, list data; Web Part properties; and user names and rights. All the data for a specific site resides in one content database. Each Web application can contain many content databases. Each site collection can be associated with only one content database, although a content database can be associated with many site collections.
- **Service application**   Service application databases store data for use by a service application. The databases for service applications vary significantly in what they are used for.

For a full list of all of the databases that support SharePoint Server 2010, see Database types and descriptions (SharePoint Server 2010).

### Working with the SQL Server databases that support SharePoint 2010 Products

The SQL Server databases that support SharePoint Server 2010 can be created either by SharePoint Server 2010, or by a database administrator. For more information, see Deploy by using DBA-created databases (SharePoint Server 2010) .

Microsoft does not support directly querying or modifying the databases that support SharePoint Server 2010, except for the Usage and Health Data Collection service application database, which can be queried directly and can have its schema added to.

The SQL Server databases that support SharePoint Server 2010 are subject to sizing limitations and to configuration recommendations that are not standard for SQL Server. For more information, see Storage and SQL Server capacity planning and configuration (SharePoint Server 2010).

# SQL Server as a data platform for business intelligence in SharePoint 2010 Products

SharePoint Server 2010 can be used with SQL Server BI tools to analyze and display BI data in meaningful ways. SQL Server provides the primary data infrastructure and business intelligence platform that gives report authors and business users trusted, scalable, and secure data.

The following sections describe the technologies and features in SQL Server that support business intelligence functionality and features in SharePoint Server 2010.

## SQL Server database engine

The SQL Server database engine is the core service for storing, processing, and securing data. BI data can be collected from the SQL Server database engine. For more information, see SQL Server Database Engine (http://go.microsoft.com/fwlink/?LinkId=199540).

## SQL Server Analysis Services (SSAS): multi-dimensional data

Microsoft SQL Server Analysis Services (SSAS) multidimensional data enables you to design, create, and manage multidimensional structures that contain detail and aggregated data from multiple data sources. A cube wizard is available in SQL Server 2008 R2 that simplifies how you can create cubes. Dimensional data or cube data is a prototypical data source for the types of analysis that can be done by using the business intelligence-related service applications in SharePoint Server 2010. To learn how relational and multi-dimensional data helps users analyze data, see Data warehousing, OLAP, and Analysis Services for SharePoint 2010. For more information, see SQL Server Analysis Services - Multidimensional Data (http://go.microsoft.com/fwlink/?LinkId=199541).

## SQL Server Analysis Services: data mining

SQL Server Analysis Services data mining tools provide a set of industry-standard data mining algorithms and other tools that help you discover trends and patterns in your data. The following Excel add-ins help you perform predictive analysis:

- Table Analysis Tools for Excel provide easy-to-use tools that take advantage of  Analysis Services Data Mining to perform powerful analytics on spreadsheet data. For more information, see SQL Server Analysis Services - Data Mining (http://go.microsoft.com/fwlink/?LinkId=199543).
- Data Mining Client for Excel lets users build, test, and query data mining models within Microsoft Office Excel 2007 by using either worksheet data or external data available through Analysis Services.

**Note:**
To enable add-ins, you must have a connection to the server.

## SQL Server Reporting Services (SSRS)

Microsoft SQL Server Reporting Services (SSRS) and SharePoint Server 2010 are easily integrated. SQL Server Reporting Services has a full range of tools with which you can create, deploy, and manage reports for your organization. It also has features that enable you to extend and customize your reporting functionality.

The available functionality includes:

- Creating reports with Report Builder 3, one of the SQL Server Reporting Services authoring tools, which you can launch directly from SharePoint Server 2010.
- Publishing SSRS reports in SharePoint Server 2010.

  You can publish report server content types to a SharePoint library and then view and manage those documents from a SharePoint site.

For more information about SSRS, see SQL Server Reporting Services (http://go.microsoft.com/fwlink/?LinkId=199545). For more information about how to install the different integration modes, see Overview of documentation for SQL Server Reporting Services reports in SharePoint.

## SQL Server Integration Services (SSIS)

Microsoft SQL Server Integration Services (SSIS) provides rich data integration and data transformation solutions. You can create a repeatable extract, transform, and load (ETL) process to automate moving data from sources such as XML data files, flat files, or relational data sources to one or more destinations. If data comes from disparate sources and is not mined or cleansed for the benefits that are provided in BI applications, SQL Server Integration Services helps prepare the data. For more information, see SQL Server Integration Services (http://go.microsoft.com/fwlink/?LinkId=199546).

## Business Intelligence Development Studio (BIDS)

Microsoft Business Intelligence Development Studio (BIDS) provides intuitive wizards for building integration, reporting, and analytic solutions in a unified environment. BIDS supports the complete development life cycle of developing, testing, and deploying solutions and reports. BIDS is based on the Visual Studio 2005 development environment but customizes it with the SQL Server services–specific extensions and project types for reports, ETL data flows, OLAP cubes, and data mining structure.

## PowerPivot for Excel and PowerPivot for SharePoint

PowerPivot is an add-in that enables users to create self-service BI solutions. It also facilitates sharing and collaboration on those solutions in a SharePoint Server 2010 environment. PowerPivot also enables IT organizations to increase operational efficiencies through Microsoft SQL Server 2008 management tools.  Components of PowerPivot include the following:

- PowerPivot for Excel 2010 is a data analysis add-in that delivers computational power directly to Microsoft Excel 2010. PowerPivot for Excel (formerly known as "Gemini") lets users analyze large quantities of data, and its integration with SharePoint Server 2010 helps IT departments monitor and manage how users collaborate. The add-in removes the one-million-row limit for worksheets and provides rapid calculations for large data sets. For more information, see PowerPivot Overview (http://go.microsoft.com/fwlink/?LinkId=199547).
- PowerPivot for SharePoint 2010 extends SharePoint Server 2010 and Excel Services to add server-side processing, collaboration, and document management support for the PowerPivot workbooks that you publish to SharePoint sites. For more information, see PowerPivot for SharePoint (http://go.microsoft.com/fwlink/?LinkId=199547).

## Master Data Services

SQL Server Master Data Services lets you centrally manage important data assets companywide and across diverse systems to provide more trusted data to your BI applications. Master Data Services helps you create a master data hub that includes a thin-client data management application for a data steward. The application can also apply workflow to assigned owners, apply extensible business rules to safeguard data quality, and apply hierarchy and attribute management strategies. For more information, see Master Data Services (http://go.microsoft.com/fwlink/?LinkId=199548).

## StreamInsight and complex event processing

Microsoft StreamInsight is a new feature in SQL Server 2008 R2 that provides a powerful platform for developing and deploying complex event processing (CEP) applications. CEP is a technology for processing streams of events with high-throughput and low-latency. StreamInsight lets you analyze data without first storing it, and helps you monitor data from multiple sources to detect patterns, trends, and exceptions almost instantly.  The ability to monitor, analyze, and act on data in motion in an event-driven manner provides significant opportunity to make more rapid, informed business decisions. For more information, see Microsoft StreamInsight (http://go.microsoft.com/fwlink/?LinkId=199549).

# SharePoint Server 2010 authoring and publishing tools for business intelligence

The following are authoring and publishing tools in SharePoint Server 2010 that contribute to how to create KPIs, scorecards, dashboards, and reports. Each of the tools can link to SQL Server relational and multi-dimensional data. To learn more, see Data warehousing, OLAP, and Analysis Services for SharePoint 2010. For more information about the services and to see how each tool uses SQL Server data, see Architecture for business intelligence in SharePoint Server 2010 and Choosing a business intelligence tool in SharePoint Server.

- **Excel Services** Use Excel 2010 and Excel Services to view, refresh, and interact with analytic models connected to data sources. Also use them for analysis, filtering, and presenting locally stored data. Excel 2010 is the authoring tool. Excel Services lets you publish Excel 2010 files to SharePoint Server 2010.
- **Visio Services** Use Visio Services to build a visual representation of business structures that are bound to data. Examples include creating visual processes, systems, and resources that show visual performance. For example, an engineer can use the visualization to create data-bound objects to represent a process.
- **PerformancePoint Services** Use PerformancePoint Services to create dashboards, scorecards, and key performance indicators (KPIs) that deliver a summarized view of business performance. PerformancePoint Services gives users integrated analytics for monitoring, analyzing, and reporting.
- **Web Analytics service application** Use the Web Analytics service application to understand more about visits to the SharePoint sites. The Web Analytics service application collects data about how end-users access SharePoint pages.

# Related content

| | |
|---|---|
| Resource center | Business Continuity Management for SharePoint Server 2010<br>Business Intelligence in SharePoint Server 2010 (http://go.microsoft.com/fwlink/?LinkId=199757)<br>Microsoft Business Intelligence (http://go.microsoft.com/fwlink/?LinkId=199758)<br>SQL Server Tech Center (http://go.microsoft.com/fwlink/?LinkId=199760)<br>SQL Server Analysis Services Multidimensional |

| | |
|---|---|
| | Data (SSAS)<br>(http://go.microsoft.com/fwlink/?LinkId=199761)<br>SQL Server Analysis Services (SSAS) Data Mining<br>(http://go.microsoft.com/fwlink/?LinkId=199762) |
| Developer content | SharePoint Developer Center<br>(http://go.microsoft.com/fwlink/?LinkID=159918)<br>SQL Server Developer Center<br>(http://go.microsoft.com/fwlink/?LinkId=199764)<br>SQL Server Database Engine<br>(http://go.microsoft.com/fwlink/?LinkId=199765)<br>SQL Server Reporting Services (SSRS)<br>(http://go.microsoft.com/fwlink/?LinkId=199766)<br>SQL Server StreamInsight<br>(http://go.microsoft.com/fwlink/?LinkId=199767) |

# SQL Server 2008 R2 and SharePoint 2010 Products: Better Together (white paper) (SharePoint Server 2010)

Choosing an edition of Microsoft SQL Server 2008 R2 is an important step when planning a Microsoft SharePoint Server 2010 deployment. This paper describes the benefits of deploying on SQL Server 2008 R2 Enterprise Edition and scenarios in which its features can be applied.

Download this white paper as a Word document (.docx) (http://go.microsoft.com/fwlink/?LinkID=187264).

# Storage and SQL Server capacity planning and configuration (SharePoint Server 2010)

This article describes how to plan for and configure the storage and Microsoft SQL Server database tier in a Microsoft SharePoint Server 2010 environment.

The capacity planning information in this document provides guidelines for you to use in your planning. It is based on testing performed at Microsoft on live properties. However, your results may vary based on the equipment you use and the features and functionality that you implement for your sites.

Because SharePoint Server 2010 often runs in environments in which databases are managed by separate SQL Server database administrators, this document is intended for joint use by SharePoint Server 2010 farm implementers and SQL Server database administrators. It assumes significant understanding of both SharePoint Server 2010 and SQL Server.

This article assumes that you are familiar with the concepts presented in Capacity management and sizing for SharePoint Server 2010.

## Design and configuration process for SharePoint 2010 Products storage and database tier

We recommend that you break the storage and database tier design process into the following steps. Each section provides detailed information about each design step, including storage requirements and best practices:

- Gather storage and SQL Server space and I/O requirements
- Choose SQL Server version and edition
- Design storage architecture based on capacity and I/O requirements
- Estimate memory requirements
- Understand network topology requirements
- Configure SQL Server
- Validate and monitor storage and SQL Server performance

## Gather storage and SQL Server space and I/O requirements

Several SharePoint Server 2010 architectural factors influence storage design. The amount of content, features and service applications used, number of farms, and availability needs are key factors.

Before you start to plan storage, you should understand the databases that SharePoint Server 2010 can use.

In this section:

- Databases used by SharePoint 2010 Products
- Understand SQL Server and IOPS
- Estimate core storage and IOPS needs
- Estimate service application storage needs and IOPS
- Determine availability needs

## Databases used by SharePoint 2010 Products

The databases that are installed with SharePoint Server 2010 depend on the features that are being used in the environment. All SharePoint 2010 Products environments rely on the SQL Server system databases. This section provides a summary of the databases installed with SharePoint Server 2010. For detailed information, see Database types and descriptions (SharePoint Server 2010) and Database model (http://go.microsoft.com/fwlink/?LinkId=187968).

| Product version and edition | Databases |
|---|---|
| SharePoint Foundation 2010 | Configuration<br>Central Administration content<br>Content (one or more)<br>Usage and Health Data Collection<br>Business Data Connectivity<br>Application Registry service (if upgrading from Microsoft SharePoint Server 2010 2007 Business Data Catalog)<br>Subscription Settings service (if it is enabled through Windows PowerShell) |
| Additional databases for SharePoint Server 2010 Standard edition | Search service application:<br>• Search administration<br>• Crawl (one or more)<br>• Properties (one or more)<br><br>User Profile service application:<br>• Profile<br>• Synchronization<br>• Social tagging |

| Product version and edition | Databases |
|---|---|
| | Web analytics service application<br>• Staging<br>• Reporting<br><br>Secure store<br>State<br>Managed Metadata<br>Word Automation services |
| Additional databases for SharePoint Server 2010 Enterprise edition | PerformancePoint |
| Additional databases for Project Server 2010 | Draft<br>Published<br>Archive<br>Reporting |
| Additional database for FAST Search Server | Search administration |

If you are integrating more fully with SQL Server, your environment may also include additional databases, as in the following scenarios:

- Microsoft SQL Server 2008 R2 PowerPivot for Microsoft SharePoint 2010 can be used in a SharePoint Server 2010 environment that includes SQL Server 2008 R2 Enterprise Edition and SQL Server Analysis Services. If in use, you must also plan to support the PowerPivot Application database, and the additional load on the system. For more information, see Plan a PowerPivot deployment in a SharePoint farm (http://go.microsoft.com/fwlink/?LinkID=186698).
- The Microsoft SQL Server 2008 Reporting Services (SSRS) plug-in can be used with any SharePoint 2010 Products environment. If you are using the plug-in, plan to support the two SQL Server 2008 Reporting Services databases and the additional load that is required for SQL Server 2008 Reporting Services.

## Understand SQL Server and IOPS

On any server that hosts SQL Server, it is very important that the server achieve the fastest response possible from the I/O subsystem.

More and faster disks or arrays provide sufficient I/O operations per second (IOPS) while maintaining low latency and queuing on all disks.

Slow response from the I/O subsystem cannot be compensated for by adding other types of resources such as CPU or memory; however, it can influence and cause issues throughout the farm. Plan for minimal latency before deployment, and monitor your existing systems.

Before you deploy a new farm, we recommend that you benchmark the I/O subsystem by using the SQLIO disk subsystem benchmark tool. For details, see SQLIO Disk Subsystem Benchmark Tool (http://go.microsoft.com/fwlink/?LinkID=105586).

For detailed information about how to analyze IOPS requirements from a SQL Server perspective, see Analyzing I/O Characteristics and Sizing Storage Systems for SQL Server Database Applications (http://sqlcat.com/whitepapers/archive/2010/05/10/analyzing-i-o-characteristics-and-sizing-storage-systems-for-sql-server-database-applications.aspx).

# Estimate core storage and IOPS needs

Configuration and content storage and IOPs are the base layer that you must plan for in every SharePoint Server 2010 deployment.

## Configuration storage and IOPS

Storage requirements for the Configuration database and the Central Administration content database are not large. We recommend that you allocate 2 GB for the Configuration database and 1 GB for the Central Administration content database. Over time, the Configuration database may grow beyond 1 GB, but it does not grow quickly — it grows by approximately 40 MB for each 50,000 site collections.

Transaction logs for the Configuration database can be large, therefore we recommend that you change the recovery model for the database from full to simple.

**Note:**

If you want to use SQL Server database mirroring to provide availability for the Configuration database, you must use the full recovery model.

IOPS requirements for the Configuration database and Central Administration content database are minimal.

## Content storage and IOPS

Estimating the storage and IOPS required for content databases is not a precise activity. In testing and explaining the following information, we intend to help you derive estimates to use for determining the initial size of your deployment. However, when your environment is running, we expect that you will revisit your capacity needs based on the data from your live environment.

For more information about our overall capacity planning methodology, see Capacity management and sizing for SharePoint Server 2010.

### Estimate content database storage

The following process describes how to approximately estimate the storage required for content databases, without considering log files:

1. Calculate the expected number of documents. This value is referred to as  in the formula.

   How you calculate the number of documents will be determined by the features that you are using. For example, for My Site Web sites or collaboration sites, we recommend that you calculate the expected number of documents per user and multiply by the number of users. For records management or content publishing sites, you may calculate the number of documents that are managed and generated by a process.

   If you are migrating from a current system, it may be easier to extrapolate your current growth rate and usage. If you are creating a new system, review your existing file shares or other repositories and estimate based on that usage rate.

2. Estimate the average size of the documents that you will be storing. This value is referred to as  in the formula.  It may be worthwhile to estimate averages for different types or groups of sites. The average file size for My Site Web sites, media repositories, and different department portals can vary significantly.
3. Estimate the number of list items in the environment. This value is referred to as  in the formula.

   List items are more difficult to estimate than documents. We generally use an estimate of three times the number of documents (), but this will vary based on how you expect to use your sites.

4. Determine the approximate number of versions. Estimate the average number of versions any document in a library will have (this value will usually be much lower than the maximum allowed number of versions). This value is referred to as  in the formula.

   The value of  must be above zero.

5. Use the following formula to estimate the size of your content databases:

   Database size = (( × ) × ) + (10 KB × ( + ( × )))

   The value of 10 KB in the formula is a constant that roughly estimates the amount of metadata required by SharePoint Server 2010. If your system requires significant use of metadata, you may want to increase this constant.

As an example, if you were to use the formula to estimate the amount of storage space required for the data files for a content database in a collaboration environment with the following characteristics, you would need approximately 105 GB.

| Input | Value |
|---|---|
| Number of documents () | 200,000<br>Calculated by assuming 10,000 users times 20 documents |
| Average size of documents () | 250 KB |
| List items () | 600,000 |
| Number of non-current versions () | 2<br>Assuming that the maximum versions allowed is |

| Input | Value |
|---|---|
| | 10 |

Database size = ((( x )) × ) + (( KB × ( + ( x ))) = KB or GB

### Features that influence the size of content databases

The use of the following SharePoint Server 2010 features can significantly affect the size of content databases:

- **Recycle bins** Until a document is fully deleted from both the first stage and second stage recycle bin, it occupies space in a content database. Calculate how many documents are deleted each month to determine the effect of recycle bins on the size of content databases. For more information, see Configure Recycle Bin settings (SharePoint Server 2010).
- **Auditing** Audit data can quickly compound and use large amounts of space in a content database, especially if view auditing is turned on. Rather than letting audit data grow without restraint, we recommend that you only enable auditing on the events that are important to meet regulatory needs or internal controls. Use the following guidelines to estimate the space you will need to reserve for auditing data:
  - Estimate the number of new auditing entries for a site, and multiply this number by 2 KB (entries generally are limited to 4 KB, with an average size of about 1 KB).
  - Based on the space that you want to allocate, determine the number of days of audit logs you want to keep.
- Office Web Apps. If Office Web Apps are being used, the Office Web Apps cache can significantly affect the size of a content database. By default, the Office Web Apps cache is configured to be 100 GB. For more information about the size of the Office Web Apps cache, see Manage the Office Web Apps cache.

### Estimate content database IOPS requirements

IOPS requirements for content databases vary significantly based on how your environment is being used, and how much disk space and how many servers you have. In general, we recommend that you compare the predicted workload in your environment to one of the solutions that we tested. For more information, see Performance and capacity test results and recommendations (SharePoint Server 2010).

**Important:**
The testing for the content in this section is not yet complete. Check back for additional information.

## Estimate service application storage needs and IOPS

After estimating content storage and IOPs needs, you must next determine the storage and IOPs required by the service applications that are being used in your environment.

### SharePoint Foundation 2010 service application storage and IOPS requirements

To estimate the storage requirements for the service applications in the system, you must first be aware of the service applications and how you will use them. Service applications that are available in SharePoint Foundation 2010 that have databases are listed in the following table.

| Service application database | Size estimation recommendation |
|---|---|
| Usage and Health Data Collection | The Usage database can grow very quickly and require significant IOPS.<br>For example, in collaborative environments that use out-of-the-box settings, 1 million HTTP requests require 2 GB of storage.<br>Use one of the following formulas to estimate the amount of IOPS required:<br>• 115 × page hits/second<br>• 5 × HTTP requests<br>If you must restrict the size of the usage database, we recommend that you start by logging only page requests. You can also restrict the size of the database by setting the default interval of data to be kept to be less than two weeks.<br>If possible, put the Usage database on its own disk or spindle. |
| Business Data Connectivity service | The size of the Business Data Connectivity services database is primarily affected by the number of external content types that you plan to support. Allocate 0.5 MB for each external content type. If you don't know how many external content types you might need, we recommend that you allocate 50 MB. IOPS requirements are minimal. |
| Application Registry service | Allocate 1 GB only if you are upgrading from the Microsoft Office SharePoint Server 2007 Business Data Catalog. IOPS requirements are minimal. |
| Subscription settings | Allocate 1 GB. IOPS requirements are minimal |

### SharePoint Server 2010 service application storage and IOPs requirements

To estimate the storage requirements for the service applications in the system, you must first be aware of the service applications and how you will use them. Service applications that are available in SharePoint Server 2010 that have databases are listed in the following table.

| Service application | Size estimation recommendation |
|---|---|
| Search | Search requires three databases. Your environment may include multiple Property and Crawl databases.<br>The Search administration database is typically small: allocate 10 GB.<br>To estimate the required storage for your Property and Crawl databases, use the following multipliers:<br>• Crawl: 0.046 × (sum of content databases)<br>• Property: 0.015 × (sum of content databases)<br>The IOPS requirements for Search are significant.<br>• For the Crawl database, search requires from 3,500 to 7,000 IOPS.<br>• For the Property database, search requires 2,000 IOPS.<br>For detailed information about how to estimate capacity required for Search, see Performance and capacity test results and recommendations (SharePoint Server 2010). |
| User Profile | The User Profile service application is associated with three databases: Profile, Synch, and Social Tagging.<br>To estimate the required storage for the databases, use the following information:<br>• Profile. With out-of-the-box settings, in an environment configured to use Active Directory, the profile database requires approximately 1 MB per user profile.<br>• Synchronization. With out-of-the-box settings, in  an environment that has few groups per user, the synch database requires approximately 630 KB per user profile. 90% of the space will be used by the data file. |

<table>
<tr><th>Service application</th><th>Size estimation recommendation</th></tr>
<tr><td></td><td>• Social tagging. With out-of-the-box settings, the social tagging database requires approximately 0.009 MB per tag, comment, or rating. To estimate how many tags and notes users will create, consider the following information about the site del.icio.us:<br>◦ Approximately 10% of users are considered active.<br>◦ Active users create 4.5 tags and 1.8 comments per month.<br>In a live collaboration environment with 160,000 user profiles, 5 groups, 79,000 tags, comments and ratings (2,500 comments, 76,000 tags, and 800 ratings), and out-of-the-box settings, we saw the following sizes for these databases:<br><table>
<tr><th>Database name</th><th>Database size</th></tr>
<tr><td>Profile</td><td>155 GB</td></tr>
<tr><td>Synchronization</td><td>96 GB</td></tr>
<tr><td>Social tagging</td><td>0.66 GB</td></tr>
</table></td></tr>
<tr><td>Managed metadata</td><td>The Managed Metadata service application has one database. The size of the database is affected by the number of content types and keywords used in the system. Many environments will include multiple instances of the Managed Metadata service application. For detailed information about how to estimate the size and IOPS requirements for this database, see Performance and capacity test results and recommendations (SharePoint Server 2010).</td></tr>
<tr><td>Web Analytics</td><td>Web Analytics has two databases: Staging and Reporting. Many factors influence the size of the databases. They include retention period, the daily volume of data being tracked, and the number of</td></tr>
</table>

| Service application | Size estimation recommendation |
|---|---|
| | site collections, sites, and subsites in the Web application being analyzed. For detailed information about how to estimate their sizing and IOPS requirements, see Performance and capacity test results and recommendations (SharePoint Server 2010). |
| Secure store | The size of the Secure Store service application database is determined by the number of credentials in the store and the number of entries in the audit table. We recommend that you allocate 5 MB for each 1,000 credentials for it. It has minimal IOPS. |
| State | The State service application has one database. We recommend that you allocate 1 GB for it. It has minimal IOPS. |
| Word Automation service | The Word Automation service application has one database. We recommend that you allocate 1 GB for it. It has minimal IOPS. |
| PerformancePoint | The PerformancePoint service application has one database. We recommend that you allocate 1 GB for it. It has minimal IOPS. |

## Determine availability needs

Availability is the degree to which a SharePoint Server 2010 environment is perceived by users to be available. An available system is a system that is resilient — that is, incidents that affect service occur infrequently, and timely and effective action is taken when they do occur.

Availability requirements can significantly increase your storage needs. For detailed information, see Plan for availability (SharePoint Server 2010).

# Choose SQL Server version and edition

Although SharePoint 2010 Products can run on Microsoft SQL Server 2008 R2, SQL Server 2008, or SQL Server 2005, we strongly recommend that you consider running your environment on the Enterprise Edition of SQL Server 2008 or SQL Server 2008 R2 to take advantage of the additional performance, availability, security, and management capabilities that it provides. For more information about the benefits of using SQL Server 2008 R2 Enterprise Edition, see SQL Server 2008 R2 and SharePoint 2010 Products: Better Together (white paper) (SharePoint Server 2010).

In particular, you should consider your need for the following features:

- **Backup compression** Backup compression can speed up any SharePoint backup, and is available in SQL Server 2008 Enterprise Edition or SQL Server 2008 R2 Standard edition. By setting the compression option in your backup script, or by configuring the server that is running SQL Server to compress by default, you can significantly reduce the size of your database backups and shipped logs. For more information, see Backup Compression (SQL Server) (http://go.microsoft.com/fwlink/?LinkId=129381&clcid=0x409).

  **Note:**
  SQL Server data compression is not supported for SharePoint 2010 Products.

- **Transparent data encryption** If your security requirements include the need for transparent data encryption, you must use SQL Server Enterprise Edition.
- **Web Analytics service application** If you plan to use the Web Analytics service application for significant analysis, consider SQL Server Enterprise Edition so that the system can take advantage of table partitioning.
- **Content deployment** If you plan to use the content deployment feature, consider SQL Server Enterprise Edition so that the system can take advantage of SQL Server database snapshots.
- **Remote BLOB storage** If you want to take advantage of remote BLOB storage to a database or location outside the files associated with each content database, you must use SQL Server 2008 or SQL Server 2008 R2 Enterprise Edition.
- **Resource governor** Resource Governor is a technology introduced in SQL Server 2008 that enables you to manage SQL Server workloads and resources by specifying limits on resource consumption by incoming requests. Resource Governor enables you to differentiate workloads and allocate CPU and memory as they are requested, based on the limits that you specify. It is available only in SQL Server 2008 or SQL Server 2008 R2 Enterprise edition. For more information about using Resource Governor, see Managing SQL Server Workloads with Resource Governor.

  We recommend that you use Resource Governor with SharePoint Server 2010 to:

  - Limit the amount of SQL Server resources that the Web servers targeted by the search crawl component consume. As a best practice, we recommend limiting the crawl component to 10 percent CPU when the system is under load.
  - Monitor how many resources are consumed by each database in the system — for example, you can use Resource Governor to help you determine the best placement of databases among computers that are running SQL Server.

- **PowerPivot for SharePoint 2010** Enables users to share and collaborate on user-generated data models and analysis in Excel and in the browser while automatically refreshing those analyses. It is part of SQL Server 2008 R2 Enterprise Edition Analysis Services.

# Design storage architecture based on capacity and I/O requirements

The storage architecture and disk types that you select for your environment can affect system performance.

In this section:

- Choose a storage architecture
- Choose disk types
- Choose RAID types

## Choose a storage architecture

Direct Attached Storage (DAS), Storage Area Network (SAN), and Network Attached Storage (NAS) storage architectures are supported with SharePoint Server 2010, although NAS is only supported for use with content databases that are configured to use remote BLOB storage. Your choice depends on factors within your business solution and your existing infrastructure.

Any storage architecture must support your availability needs and perform adequately in IOPS and latency. To be supported, the system must consistently return the first byte of data within 20 milliseconds (ms).

### Direct Attached Storage (DAS)

DAS is a digital storage system that is directly attached to a server or workstation, without a storage network in between. DAS physical disk types include Serial Attached SCSI (SAS) and Serial Attached ATA (SATA).

In general, we recommend that you choose a DAS architecture when a shared storage platform cannot guarantee a response time of 20 ms and sufficient capacity for average and peak IOPs.

### Storage Area Network (SAN)

SAN is an architecture to attach remote computer storage devices (such as disk arrays and tape libraries) to servers in such a way that the devices appear as locally attached to the operating system (for example, block storage).

In general, we recommend that you choose a SAN when the benefits of shared storage are important to your organization.

The benefits of shared storage include the following:

- Easier to reallocate disk storage between servers.
- Can serve multiple servers.
- No limitations on the number of disks that can be accessed.

### Network Attached Storage (NAS)

A NAS unit is a self-contained computer that is connected to a network. Its sole purpose is to supply file-based data storage services to other devices on the network. The operating system and other software on the NAS unit provide the functionality of data storage, file systems, and access to files, and the management of these functionalities (for example, file storage).

**Note:**
NAS is only supported for use with content databases that are configured to use remote BLOB storage. Any network storage architecture must respond to a ping within 1 ms and must return the first byte of data within 20 ms. This restriction does not apply to the local SQL Server FILESTREAM provider, because it only stores data locally on the same server.

## Choose disk types

The disk types that you use in the system can affect reliability and performance. All else being equal, larger drives increase mean seek time. SharePoint Server 2010 supports the following types of drives:

- Small Computer System Interface (SCSI)
- Serial Advanced Technology Attachment (SATA)
- Serial-attached SCSI (SAS)
- Fibre Channel (FC)
- Integrated Device Electronics (IDE)
- Solid State Drive (SSD) or Flash Disk

## Choose RAID types

RAID (Redundant Array of Independent Disks) is often used to both improve the performance characteristics of individual disks (by striping data across several disks) and to provide protection from individual disk failures.

All RAID types are supported for SharePoint Server 2010; however, we recommend that you use RAID 10 or a vendor-specific RAID solution that has equivalent performance.

When you configure a RAID array, make sure that you align the file system to the offset that is supplied by the vendor. In the absence of vendor guidance, refer to SQL Server Predeployment I/O Best Practices (http://go.microsoft.com/fwlink/?LinkID=105583).

For more information about provisioning RAID and the SQL Server I/O subsystem, see SQL Server Best Practices Article (http://go.microsoft.com/fwlink/?LinkId=168612).

# Estimate memory requirements

The memory required for SharePoint Server 2010 is directly related to the size of the content databases that you are hosting on a server that is running SQL Server.

As you add service applications and features, your requirements are likely to increase. The following table gives guidelines for the amount of memory we recommend.

**Note:**
Our definitions of small and medium deployments are those described in the "Reference Architectures" section of the article Capacity management and sizing for SharePoint Server 2010.

| **Combined size of content databases** | **RAM recommended for computer running SQL Server** |
|---|---|
| Minimum for small production deployments | 8 GB |
| Minimum for medium production deployments | 16 GB |
| Recommendation for up to 2 terabytes | 32 GB |
| Recommendation for the range of 2 terabytes to a maximum of 5 terabytes | 64 GB |

**Note:**
These values are higher than those recommended as the minimum values for SQL Server because of the distribution of data required for a SharePoint Server 2010 environment. For more information about SQL Server system requirements, see Hardware and Software Requirements for Installing SQL Server 2008 (http://go.microsoft.com/fwlink/?LinkId=129377).

Other factors that may influence the memory required include the following:

- The use of SQL Server mirroring.
- The frequent use of files larger than 15 megabytes (MB).

# Understand network topology requirements

Plan the network connections within and between farms. We recommend that you use a network that has low latency.

The following list provides some best practices and recommendations:

- All servers in the farm should have LAN bandwidth and latency to the server that is running SQL Server. Latency should be no greater than 1 ms.
- We do not recommend a wide area network (WAN) topology in which a server that is running SQL Server is deployed remotely from other components of the farm over a network that has latency greater than 1 ms. This topology has not been tested.

- Plan for an adequate WAN network if you are planning to use SQL Server mirroring or log shipping to keep a remote site up-to-date.
- We recommend that Web servers and application servers have two network adapters: one network adapter to handle end user traffic and the other to handle communication with the servers running SQL Server.

# Configure SQL Server

The following sections describe how to plan to configure SQL Server for SharePoint Server 2010.

In this section:

- Determine how many instances or servers are required
- Configure storage and memory
- Set SQL Server options
- Configure databases

## Estimate how many servers are required

In general, SharePoint Server 2010 was designed to take advantage of SQL Server scale out — that is, SharePoint Server 2010 may perform better with a large number of medium-size servers that are running SQL Server than with only a few large servers.

Always put SQL Server on a dedicated server that is not running any other farm roles or hosting databases for any other application, unless you are deploying the system on a stand-alone server.

The following is general guidance for when to deploy an additional server that will run SQL Server:

- Add an additional database server when you have more than four Web servers that are running at full capacity.
- Add an additional database server when your content databases exceed 5 terabytes.

**Note:**
Microsoft supports server configurations that do not follow this guidance.

To promote secure credential storage when you are running the Secure Store service application, we recommend that the secure store database be hosted on a separate database instance where access is limited to one administrator.

## Configure storage and memory

On the server that is running SQL Server 2008, we recommend that the L2 cache per CPU have a minimum of 2 MB to improve memory.

### Follow vendor storage configuration recommendations

For optimal performance when you configure a physical storage array, adhere to the hardware configuration recommendations supplied by the storage vendor instead of relying on the default values of the operating system.

If you do not have guidance from your vendor, we recommend that you use the DiskPart.exe disk configuration utility to configure storage for SQL Server 2008. For more information, see Predeployment I/O Best Practices (http://go.microsoft.com/fwlink/?LinkID=105583&clcid=0x409).

### Provide as many resources as possible

Ensure that the SQL Server I/O channels to the disks are not shared by other applications, such as the paging file and Internet Information Services (IIS) logs.

Provide as much bus bandwidth as possible. Greater bus bandwidth helps improve reliability and performance. Consider that the disk is not the only user of bus bandwidth — for example, you must also account for network access.

## Set SQL Server options

The following SQL Server settings and options should be configured before you deploy SharePoint Server 2010.

- Do not enable auto-create statistics on a SQL Server that is supporting SharePoint Server 2010. SharePoint Server 2010 implements specific statistics, and no additional statistics are needed. Auto-create statistics can significantly change the execution plan of a query from one instance of SQL Server to another instance of SQL Server. Therefore, to provide consistent support for all customers, SharePoint Server 2010 provides coded hints for queries as needed to provide the best performance across all scenarios.
- To ensure optimal performance, we strongly recommend that you set **max degree of parallelism** to 1 for database servers that host SharePoint Server 2010 databases. For more information about how to set **max degree of parallelism**, see max degree of parallelism Option (http://go.microsoft.com/fwlink/?LinkId=189030).
- To improve ease of maintenance, configure SQL Server connection aliases for each database server in your farm. A connection alias is an alternative name that can be used to connect to an instance of SQL Server. For more information, see How to: Set a SQL Server Alias (SQL Server Management Studio) (http://go.microsoft.com/fwlink/?LinkId=132064&clcid=0x409).

## Configure databases

The following guidance describes best practices to plan for as you configure each database in your environment.

## Separate and prioritize your data among disks

Ideally, you should place the tempdb database, content databases, Usage database, search databases, and SQL Server 2008 transaction logs on separate physical hard disks.

The following list provides some best practices and recommendations for prioritizing data:

- When you prioritize data among faster disks, use the following ranking:

    a. Tempdb data files and transaction logs

    b. Database transaction log files

    c. Search databases, except for the Search administration database

    d. Database data files

  In a heavily read-oriented portal site, prioritize data over logs.

- Testing and customer data show that SharePoint Server 2010 farm performance can be significantly impeded by insufficient disk I/O for tempdb. To avoid this issue, allocate dedicated disks for tempdb. If a high workload is projected or monitored — that is, the average read operation or the average write operation requires more than 20 ms — you might have to ease the bottleneck by either separating the files across disks or by replacing the disks with faster disks.
- For best performance, place the tempdb on a RAID 10 array. The number of tempdb data files should equal the number of core CPUs, and the tempdb data files should be set at an equal size. Count dual core processors as two CPUs for this purpose. Count each processor that supports hyper-threading as a single CPU. For more information, see Optimizing tempdb Performance (http://go.microsoft.com/fwlink/?LinkID=148537).
- Separate database data and transaction log files across different disks. If files must share disks because the files are too small to warrant a whole disk or stripe, or you have a shortage of disk space, put files that have different usage patterns on the same disk to minimize simultaneous access requests.
- Consult your storage hardware vendor for information about how to configure all logs and the search databases for write optimization for your particular storage solution.

## Use multiple data files for content databases

Follow these recommendations for best performance:

- Only create files in the primary filegroup for the database.
- Distribute the files across separate disks.
- The number of data files should be less than or equal to the number of core CPUs. Count dual core processors as two CPUs for this purpose. Count each processor that supports hyper-threading as a single CPU.
- Create data files of equal size.

**Important:**
Although you can use the backup and recovery tools that are built in to SharePoint Server 2010 to back up and recover multiple data files, if you overwrite in the same location, the tools cannot

restore multiple data files to a different location. For this reason, we strongly recommend that when you use multiple data files for a content database, you use SQL Server backup and recovery tools. For more information about how to back up and recover SharePoint Server 2010, see Plan for backup and recovery (SharePoint Server 2010).

For more information about how to create and manage filegroups, see Physical Database Files and Filegroups (http://go.microsoft.com/fwlink/?LinkId=117909).

## Limit content database size to improve manageability

Plan for database sizing that will improve manageability, performance, and ease of upgrade for your environment.

To help ensure system performance, we strongly recommended that you limit the size of content databases to 200 GB.

A site collection should not exceed 100 GB unless it is the only site collection in the database. This limit exists so that you can use the SharePoint Server 2010 granular backup tools to move a site collection to another database if you need to.

**Important:**
Content database sizes up to 1 terabyte are supported only for large, single-site repositories and archives in which data remains reasonably static, such as reference document management systems and Records Center sites. Larger database sizes are supported for these scenarios because their I/O patterns and typical data structure formats have been designed for and tested at larger scales.

If your design requires a database larger than the recommended standard, follow this guidance:

- For databases that contain many large files that are stored as binary large objects (BLOBs), consider using remote BLOB storage (RBS). RBS is appropriate in the following circumstances:

  a. When you are running sites that contain large files that are infrequently accessed, such as knowledge repositories.

  b. When you have terabytes of data.

  c. For video or media files.

  For more information, see Plan for Remote BLOB Storage (RBS) (SharePoint Server 2010).

- Follow best practices for viewing data from large databases. For more information, see SharePoint Server 2010 capacity management: Software boundaries and limits.

For more information about large-scale document repositories, see "Estimate Performance and Capacity Requirements for Large Scale Document Repositories", available from Performance and capacity test results and recommendations (SharePoint Server 2010).

## Proactively manage the growth of data and log files

We recommend that you proactively manage the growth of data and log files by considering the following recommendations:

- As much as possible, pre-grow all data and log files to their anticipated final size.
- We recommend that you enable autogrowth for safety reasons. Do not rely on the default autogrowth settings. Consider the following guidelines when configuring autogrowth:
  - When you plan content databases that exceed the recommended size (200 GB), set the database autogrowth value to a fixed number of megabytes instead of to a percentage. This will reduce the frequency with which SQL Server increases the size of a file. Increasing file size is a blocking operation that involves filling the new space with empty pages.
  - Set the autogrowth value for the Search service application Property Store database to 10 percent.
  - If the calculated size of the content database is not expected to reach the recommended maximum size of 200 GB within the next year, set it to the maximum size the database is predicted to reach within a year — with 20 percent additional margin for error — by using the **ALTER DATABASE MAXSIZE** property. Periodically review this setting to make sure it is still an appropriate value based on past growth rates.
- Maintain a level of at least 25 percent available space across disks to allow for growth and peak usage patterns. If you are managing growth by adding disks to a RAID array or allocating more storage, monitor disk size closely to avoid running out of space.

# Validate and monitor storage and SQL Server performance

Test that your performance and backup solution on your hardware enables you to meet your service level agreements (SLAs). In particular, test the I/O subsystem of the computer that is running SQL Server to ensure that performance is satisfactory.

Test the backup solution that you are using to ensure that it can back up the system within the available maintenance window. If the backup solution cannot meet the SLAs your business requires, consider using an incremental backup solution such as System Center Data Protection Manager (DPM) 2010.

It is important to track the following resource components of a server that is running SQL Server: CPU, memory, cache/hit ratio, and I/O subsystem. When one or more of the components seems slow or overburdened, analyze the appropriate strategy based on the current and projected workload. For more information, see Troubleshooting Performance Problems in SQL Server 2008 (http://go.microsoft.com/fwlink/?LinkID=168448).

The following section lists the performance counters that we recommend that you use to monitor the performance of the SQL Server databases that are running in your SharePoint Server 2010 environment. Also listed are approximate healthy values for each counter.

For details about how to monitor performance and use performance counters, see Monitoring Performance (http://go.microsoft.com/fwlink/?LinkId=189032).

## SQL Server counters to monitor

Monitor the following SQL Server counters to ensure the health of your servers:

- **General statistics** This object provides counters to monitor general server-wide activity, such as the number of current connections and the number of users connecting and disconnecting per second from computers running an instance of SQL Server. Consider monitoring the following counter:
    - **User connections** This counter shows the amount of user connections on your computer running SQL Server. If you see this number rise by 500 percent from your baseline, you may see a performance reduction.
- **Databases** This object provides counters to monitor bulk copy operations, backup and restore throughput, and transaction log activities. Monitor transactions and the transaction log to determine how much user activity is occurring in the database and how full the transaction log is becoming. The amount of user activity can determine the performance of the database and affect log size, locking, and replication. Monitoring low-level log activity to gauge user activity and resource usage can help you to identify performance bottlenecks. Consider monitoring the following counter:
    - **Transactions/sec** This counter shows the amount of transactions on a given database or on the entire server per second. This number is more for your baseline and to help you troubleshoot issues.
- **Locks** This object provides information about SQL Server locks on individual resource types. Consider monitoring the following counters:
    - **Average Wait Time (ms)** This counter shows the average amount of wait time for each lock request that resulted in a wait.
    - **Lock Wait Time (ms)** This counter shows the wait time for locks in the last second.
    - **Lock waits/sec** This counter shows the number of locks per second that could not be satisfied immediately and had to wait for resources.
    - **Number of deadlocks/sec** This counter shows the number of deadlocks on the computer running SQL Server per second. This should not rise above 0.
- **Latches** This object provides counters to monitor internal SQL Server resource locks called latches. Monitoring the latches to determine user activity and resource usage can help you to identify performance bottlenecks. Consider monitoring the following counters:
    - **Average Latch Wait Time (ms)** This counter shows the average latch wait time for latch requests that had to wait.
    - **Latch Waits/sec** This counter shows the number of latch requests that could not be granted immediately.
- **SQL Statistics** This object provides counters to monitor compilation and the type of requests sent to an instance of SQL Server. Monitoring the number of query compilations and recompilations and the number of batches received by an instance of SQL Server gives you an indication of how quickly SQL Server is processing user queries and how effectively the query optimizer is processing the queries. Consider monitoring the following counters:

  - **SQL Compilations/sec** This counter indicates the number of times the compile code path is entered per second.
  - **SQL Re-Compilations/sec** This counter indicates the number statement recompiles per second.
- **Buffer Manager** This object provides counters to monitor how SQL Server uses memory to store data pages, internal data structures, and the procedure cache, as well as counters to monitor the physical I/O as SQL Server reads and writes database pages. Consider monitoring the following counter:
  - **Buffer Cache Hit Ratio**
  - This counter shows the percentage of pages that were found in the buffer cache without having to read from disk. The ratio is the total number of cache hits divided by the total number of cache lookups over the last few thousand page accesses. Because reading from the cache is much less expensive than reading from disk, you want this ratio to be high. Generally, you can increase the buffer cache hit ratio by increasing the amount of memory available to SQL Server.
- **Plan Cache** This object provides counters to monitor how SQL Server uses memory to store objects such as stored procedures, ad hoc and prepared Transact-SQL statements, and triggers. Consider monitoring the following counter:
  - **Cache Hit Ratio**
  - This counter indicates the ratio between cache hits and lookups for plans.

## Physical server counters to monitor

Monitor the following counters to ensure the health of your computers running SQL Server:

- **Processor: % Processor Time: _Total** This counter shows the percentage of time that the processor is executing application or operating system processes other than Idle. On the computer that is running SQL Server, this counter should be kept between 50 percent and 75 percent. In case of constant overloading, investigate whether there is abnormal process activity or if the server needs additional CPUs.
- **System: Processor Queue Length** This counter shows the number of threads in the processor queue. Monitor this counter to ensure that it remains less than two times the number of core CPUs.
- **Memory: Available Mbytes** This counter shows the amount of physical memory, in megabytes, available to processes running on the computer. Monitor this counter to ensure that you maintain a level of at least 20 percent of the total available physical RAM.
- **Memory: Pages/sec** This counter shows the rate at which pages are read from or written to disk to resolve hard page faults. Monitor this counter to ensure that it remains under 100.

For more information and memory troubleshooting methods, see SQL Server 2005 Monitoring Memory Usage (http://go.microsoft.com/fwlink/?LinkID=105585).

## Disk counters to monitor

Monitor the following counters to ensure the health of disks. Note that the following values represent values measured over time — not values that occur during a sudden spike and not values that are based on a single measurement.

- **Physical Disk: % Disk Time: DataDrive** This counter shows the percentage of elapsed time that the selected disk drive is busy servicing read or write requests–it is a general indicator of how busy the disk is. If the **PhysicalDisk: % Disk Time** counter is high (more than 90 percent), check the **PhysicalDisk: Current Disk Queue Length** counter to see how many system requests are waiting for disk access. The number of waiting I/O requests should be sustained at no more than 1.5 to 2 times the number of spindles that make up the physical disk.
- **Logical Disk: Disk Transfers/sec** This counter shows the rate at which read and write operations are performed on the disk. Use this counter to monitor growth trends and forecast appropriately.
- **Logical Disk: Disk Read Bytes/sec** and **Logical Disk: Disk Write Bytes/sec** These counters show the rate at which bytes are transferred from the disk during read or write operations.
- **Logical Disk: Avg. Disk Bytes/Read** This counter shows the average number of bytes transferred from the disk during read operations. This value can reflect disk latency — larger read operations can result in slightly increased latency.
- **Logical Disk: Avg. Disk Bytes/Write** This counter shows the average number of bytes transferred to the disk during write operations. This value can reflect disk latency — larger write operations can result in slightly increased latency.
- **Logical Disk: Current Disk Queue Length** This counter shows the number of requests outstanding on the disk at the time that the performance data is collected. For this counter, lower values are better. Values greater than 2 per disk may indicate a bottleneck and should be investigated. This means that a value of up to 8 may be acceptable for a logical unit (LUN) made up of 4 disks. Bottlenecks can create a backlog that can spread beyond the current server that is accessing the disk and result in long wait times for users. Possible solutions to a bottleneck are to add more disks to the RAID array, replace existing disks with faster disks, or move some data to other disks.
- **Logical Disk: Avg. Disk Queue Length** This counter shows the average number of both read and write requests that were queued for the selected disk during the sample interval. The rule is that there should be two or fewer outstanding read and write requests per spindle, but this can be difficult to measure because of storage virtualization and differences in RAID levels between configurations. Look for larger than average disk queue lengths in combination with larger than average disk latencies. This combination can indicate that the storage array cache is being overused or that spindle sharing with other applications is affecting performance.
- **Logical Disk: Avg. Disk sec/Read** and **Logical Disk: Avg. Disk sec/Write** These counters show the average time, in seconds, of a read or write operation to the disk. Monitor these counters to ensure that they remain below 85 percent of the disk capacity. Disk access time increases exponentially if read or write operations are more than 85 percent of disk capacity. To determine the specific capacity for your hardware, refer to the vendor documentation or use the SQLIO Disk

Subsystem Benchmark Tool to calculate it. For more information, see SQLIO Disk Subsystem Benchmark Tool (http://go.microsoft.com/fwlink/?LinkID=105586).

- **Logical Disk: Avg. Disk sec/Read** This counter shows the average time, in seconds, of a read operation from the disk. On a well-tuned system, ideal values are from 1 through 5 ms for logs (ideally 1 ms on a cached array), and from 4 through 20 ms for data (ideally less than 10 ms). Higher latencies can occur during peak times, but if high values occur regularly, you should investigate the cause.
- **Logical Disk: Avg. Disk sec/Write** This counter shows the average time, in seconds, of a write operation to the disk. On a well-tuned system, ideal values are from 1 through 5 ms for logs (ideally 1 ms on a cached array), and from 4 through 20 ms for data (ideally less than 10 ms). Higher latencies can occur during peak times, but if high values occur regularly, you should investigate the cause.

When you are using RAID configurations with the **Avg. Disk sec/Read** or **Avg. Disk sec/Write** counters, use the formulas listed in the following table to determine the rate of input and output on the disk.

| RAID level | Formula |
|---|---|
| RAID 0 | I/Os per disk = (reads + writes) / number of disks |
| RAID 1 | I/Os per disk = [reads + (2 × writes)] / 2 |
| RAID 5 | I/Os per disk = [reads + (4 × writes)] / number of disks |
| RAID 10 | I/Os per disk = [reads + (2 × writes)] / number of disks |

For example, if you have a RAID 1 system that has two physical disks, and your counters are at the values that are shown in the following table:

| Counter | Value |
|---|---|
| **Avg. Disk sec/Read** | 80 |
| **Logical Disk: Avg. Disk sec/Write** | 70 |
| **Avg. Disk Queue Length** | 5 |

The I/O value per disk can be calculated as follows: (80 + (2 × 70))/2 = 110

The disk queue length can be calculated as follows: 5/2 = 2.5

In this situation, you have a borderline I/O bottleneck.

## Other monitoring tools

You can also monitor disk latency and analyze trends by using the sys.dm_io_virtual_file_stats dynamic management view in SQL Server 2008. For more information, see sys.dm_io_virtual_file_stats (Transact-SQL) (http://go.microsoft.com/fwlink/?LinkID=105587).

# Overview of Remote BLOB Storage (SharePoint Server 2010)

This article describes how you can use Microsoft SharePoint Server 2010 together with Remote BLOB Storage (RBS) and Microsoft SQL Server 2008 Express and Microsoft SQL Server 2008 R2 Express to optimize database storage resources.

Before you implement RBS, we highly recommend that you evaluate its potential costs and benefits. For more information and recommendations about using RBS in a SharePoint Server 2010 installation, see Plan for Remote BLOB Storage (RBS) (SharePoint Server 2010).

In this article:

- Introduction to RBS
- Using RBS together with SharePoint 2010 Products

# Introduction to RBS

RBS is a library API set that is incorporated as an add-on feature pack for Microsoft SQL Server.  It can be run on the local server running Microsoft SQL Server 2008 R2, SQL Server 2008 or SQL Server 2008 R2 Express. To run RBS on a remote server, you must be running SQL Server 2008 R2 Enterprise edition. RBS is not supported for Microsoft SQL Server 2005.

Binary large objects (BLOBs) are data elements that have either of the following characteristics:

- Unstructured data that has no schema (such as a piece of encrypted data).
- A large amount of binary data (many megabytes or gigabytes) that has a very simple schema, such as image files, streaming video, or sound clips.

By default, Microsoft SQL Server stores BLOB data in its databases. As a database's usage increases, the total size of its BLOB data can expand quickly and grow larger than the total size of the document metadata and other structured data that is stored in the database. Because BLOB data can consume a lot of file space and uses server resources that are optimized for database access patterns, it can be helpful to move BLOB data out of the SQL Server database, and into a separate file.

Before RBS was supported in SQL Server, expensive storage such as RAID 10 was required for the whole SQL database including BLOB data. By using RBS, you can move 80 to 90 percent of the data (that is, BLOBs) onto less expensive storage such as RAID 5 or external storage solutions.

RBS uses a *provider* to connect to any dedicated BLOB store that uses the RBS APIs. Storage solution vendors can implement providers that work with RBS APIs. SharePoint Server 2010 supports a BLOB storage implementation that accesses BLOB data by using the RBS APIs through such a provider. You can implement RBS for Microsoft SharePoint 2010 Products by using a supported provider that you obtain from a third-party vendor. Most third-party providers store BLOBs remotely.

In addition to third-party providers, you can use the RBS FILESTREAM provider that is available through the SQL Server Remote BLOB Store installation package from the Feature Pack for Microsoft SQL Server 2008 R2. The RBS FILESTREAM provider uses the SQL Server FILESTREAM feature to store BLOBs in an additional resource that is attached to the same database and stored locally on the server. The FILESTREAM feature manages BLOBs in a SQL database by using the underlying NTFS file system.

The location that an RBS provider stores the BLOB data depends on the provider that you use. In the case of the SQL FILESTREAM provider, the data is not stored in the MDF file, but in another file that is associated with the database.

This implementation of the FILESTREAM provider is known as the *local FILESTREAM provider*. You can conserve resources by using the local RBS FILESTREAM provider to place the extracted BLOB data on a different (cheaper) local disk such as RAID 5 instead of RAID 10. You cannot use RBS with the local FILESTREAM provider on remote storage devices, such as network attached storage (NAS). The FILESTREAM provider is supported when it is used on local hard disk drives only.

A remote RBS FILESTREAM provider that is available in SQL Server 2008 R2 Express can store BLOB data on remote commodity storage such as direct-attached storage (DAS) or NAS. However, SharePoint Server 2010 does not currently support the remote RBS FILESTREAM provider.

# Using RBS together with SharePoint 2010 Products

SharePoint Server 2010 supports the FILESTREAM provider that is included in the SQL Server Remote BLOB Store installation package from the Feature Pack for SQL Server 2008 R2. This version of RBS is available at http://go.microsoft.com/fwlink/?LinkID=177388. Be aware that this is the only version of RBS that is supported by SharePoint Server 2010. Earlier versions are not supported. Third-party RBS providers can also be used with the RBS APIs to create a BLOB storage solution that is compatible with SharePoint Server 2010.

In SharePoint Server 2010, site collection backup and restore and site import or export will download the file contents and upload them back to the server regardless of which RBS provider is being used. However, the FILESTREAM provider is the only provider that is currently supported for SharePoint 2010 Products farm database backup and restore operations.

When RBS is implemented, SQL Server itself is regarded as an RBS provider. You will encounter this factor when you migrate content into and out of RBS.

If you plan to store BLOB data in an RBS store that differs from your SharePoint Server 2010 content databases, you must run SQL Server 2008 with SP1 and Cumulative Update 2. This is true for all RBS providers.

The FILESTREAM provider that is recommended for upgrading from stand-alone installations of Windows SharePoint Services 3.0 that have content databases that are over 4 gigabytes (GB) to SharePoint Server 2010 associates data locally with the current content database, and does not require SQL Server Enterprise Edition.

**Important:**

RBS does not enable any kind of direct access to any files that are stored in Microsoft SharePoint 2010 Products. All access must occur by using SharePoint 2010 Products only.

**See Also**

FILESTREAM Overview

FILESTREAM Storage in SQL Server 2008

Remote BLOB Store Provider Library Implementation Specification

# Plan for Remote BLOB Storage (RBS) (SharePoint Server 2010)

By default, SQL Server stores binary large object (BLOB) data in its databases. As a database's usage increases, the total size of the BLOB data that is stored in it can expand quickly and grow larger than the total size of the document metadata and other structured data that is stored in the database. BLOB data consumes large amounts of file space and uses server resources that are optimized for database access patterns instead of for the storage of large files.

Remote BLOB Storage (RBS) is a library API set that is incorporated as an add-on feature pack for Microsoft SQL Server.  It can be run on the local server running Microsoft SQL Server 2008 R2, SQL Server 2008 or SQL Server 2008 R2 Express. To run RBS on a remote server, you must be running SQL Server 2008 R2 Enterprise edition. RBS is designed to move the storage of BLOBs from database servers to commodity storage solutions. RBS saves significant space, conserves expensive server resources, and provides a standardized model for applications to access BLOB data. In Microsoft SharePoint Server 2010, RBS can be used for content databases only.

For more background information about RBS, including a discussion about the FILESTREAM provider, see Overview of Remote BLOB Storage (SharePoint Server 2010).

RBS can provide the following benefits:

- BLOB data can be stored on less expensive storage devices that are configured to handle simple storage.
- The administration of the BLOB storage is controlled by a system that is designed specifically to work with BLOB data.
- Database server resources are freed for database operations.

These benefits are not free. Before you implement RBS with SharePoint Server 2010, you should evaluate whether these potential benefits override the costs and limitations of implementing and maintaining RBS. This article describes this evaluation process.

In this article:

- Review the environment
- Evaluate provider options

## Review the environment

To start your analysis of RBS, review the size of the content databases. If the content database sizes meet the criteria for a RBS recommendation, you should then consider what kind of content is being accessed and how it is being used.

## Content database sizes

You can expect to benefit from RBS in the following cases:

- The content databases are larger than 500 gigabytes (GB).
- The BLOB data files are larger than 256 kilobytes (KB).
- The BLOB data files are at least 80 KB and the database server is a performance bottleneck. In this case, RBS reduces the both the I/O and processing load on the database server.

Although the presence of many small BLOBs can create some decrease in performance, the cost of storage is usually the most important consideration when you evaluate RBS. The predicted decrease in performance is usually an acceptable trade-off for the cost savings in storage hardware.

## Content type and usage

RBS is most beneficial in systems that store very large files, such as digital media. RBS is typically implemented in environments in which large stored files are infrequently accessed, such as an archive. If this situation describes your environment, you should consider implementing RBS.

If you are storing many small (less than 256 KB) files that are frequently accessed by many users, you might experience increased latency on sites that have many small files that are stored in RBS. Increased latency is one cost factor that you should consider when you evaluate RBS for your storage solution. However, it is unlikely to be the strongest consideration. The amount of increased latency is also related to the RBS provider that you use.

# Evaluate provider options

RBS requires a provider that connects the RBS APIs and SQL Server.

> **Important:**
> RBS can be run on the local server running Microsoft SQL Server 2008 R2, SQL Server 2008 or SQL Server 2008 R2 Express. To run RBS on a remote server, you must be running SQL Server 2008 R2 Enterprise edition. SharePoint Server 2010 requires you to use the version of RBS that is included with the SQL Server Remote BLOB Store installation package from the Feature Pack for Microsoft SQL Server 2008 R2. Earlier versions of RBS will not work with SharePoint Server 2010. In addition, RBS is not supported in SQL Server 2005.

BLOBs can be kept on commodity storage such as direct-attached storage (DAS) or network attached storage (NAS), as supported by the provider. The FILESTREAM provider is supported by SharePoint Server 2010 when it is used on local hard disk drives only. You cannot use RBS with FILESTREAM on remote storage devices, such as NAS.

The following table summarizes FILESTREAM benefits and limitations.

| Operational requirement | RBS with FILESTREAM | RBS without FILESTREAM |
|---|---|---|
| SQL Server integrated backup and recovery of the BLOB Store | Yes | Yes |
| Scripted migration to BLOBs | Yes | Yes |
| Supports mirroring | No | No |
| Log shipping | Yes | Yes, with provider implementation |
| Database snapshots | No[1] | No[1] |
| Geo replication | Yes | No |
| Encryption | NTFS only | No |
| Network Attached Storage (NAS) | Not supported by SharePoint 2010 Products | Yes, with provider implementation |

[1] If the RBS provider that you are using does not support snapshots, you cannot use snapshots for content deployment or backup. For example, the SQL FILESTREAM provider does not support snapshots.

If FILESTREAM is not a practical provider for your environment, you can purchase a supported third-party provider. In this case, you should evaluate the following criteria when shopping for a provider:

- Backup and restore capability
- Tested disaster recovery
- Deployment and data migration
- Performance impact
- Long-term administrative costs

**Important:**
We do not recommend that you develop your own provider unless you are an independent software vendor (ISV) that has significant development experience in designing storage solutions.

# Plan for business continuity management (SharePoint Server 2010)

Business continuity management consists of the business decisions, processes, and tools you put in place in advance to handle crises. A crisis might affect your business only, or be part of a local, regional, or national event.

Features of Microsoft SharePoint Server 2010 are likely to be part of your business continuity management strategy, but your overall plan should be much more comprehensive and include the following elements:

- Clearly documented procedures.
- Offsite storage of key business records.
- Clearly designated contacts.
- Ongoing staff training, including practices and drills.
- Offsite recovery mechanisms.

In this article:

- Business continuity management capabilities
- Service level agreements

## Business continuity management capabilities

Microsoft SharePoint Server 2010 includes the following capabilities that support business continuity management.

- **Versioning** Users can lose data by overwriting a document. With versioning, users can keep multiple versions of the same document in a document library. In the event of an unwanted change, an overwritten document, or document corruption, the previous version can easily be restored by the user. When versioning is enabled, users can recover their data themselves.

  For more information, see Plan to protect content by using recycle bins and versioning (SharePoint Server 2010).

- **Recycle Bin** SharePoint Server 2010 includes a two-stage Recycle Bin. Users who have the appropriate permissions can use the first-stage Recycle Bin to recover documents, list items, lists, and document libraries that have been deleted from a site. Site collection administrators can use the second-stage Recycle Bin, also called the Site Collection Recycle Bin, to recover items that have been deleted from the first-stage Recycle Bin. When the first-stage Recycle Bin is enabled, users can recover their data themselves.

  For more information, see Plan to protect content by using recycle bins and versioning (SharePoint Server 2010).

- **Records Center** Records Center sites support managing records storage for legal, regulatory, or business reasons. For more information, see Records management planning (SharePoint Server 2010).
- **Backup and recovery** You can use Windows PowerShell cmdlets or the SharePoint Central Administration Web site to back up and recover farms, databases, Web applications, and site collections. There are also many external and third-party tools that you can use to back up and recover data. For more information, see Plan for backup and recovery (SharePoint Server 2010).
- **Availability** No single feature provides availability within a SharePoint Server 2010 environment. You can choose among many approaches to improve availability, including the following:
  - Fault tolerance of components and the network.
  - Redundancy of server roles and servers within a farm.

  For more information about availability, see Plan for availability (SharePoint Server 2010).
- **Disaster recovery** No single feature provides disaster recovery within a SharePoint Server 2010 environment. You can choose among many approaches to improve availability when a data center goes offline, including the following:
  - Offsite storage of backups, both within and outside your region.
  - Shipping images of servers to offsite locations.
  - Running multiple data centers, but serving data only through one, keeping the others available on standby.

  For more information about disaster recovery, see Plan for disaster recovery (SharePoint Server 2010).

# Service level agreements

Business continuity management is a key area in which IT groups offer service level agreements (SLAs) to set expectations with customer groups. Many IT organizations offer various SLAs that are associated with different chargeback levels.

The following list describes common features of business continuity management SLAs:

- Versioning
  - Whether offered.
  - Amount of space allocated.
- Recycle Bins
  - Whether offered.
  - Amount of space allocated for the first-stage Recycle Bin and second-stage Recycle Bin.
  - Length of time that items are held before they are permanently deleted in each Recycle Bin.
  - Additional charges for recovering items that have been permanently deleted from the second-stage Recycle Bin.
- Backup and recovery

  Backup and recovery SLAs usually identify objects and services that can be backed up and recovered, and the recovery time objective, recovery point objective, and recovery level objective for each. The SLA may also identify the available backup window for each object. For more information about backup and recovery SLAs, see Plan for backup and recovery (SharePoint Server 2010).

  - *Recovery time objective (RTO)* is the objective for the maximum time a data recovery process will take. It is determined by the amount of time the business can afford for the site or service to be unavailable.
  - *Recovery point objective (RPO)* is the objective for the maximum amount of time between the last available backup and any potential failure point. It is determined by how much data the business can afford to lose in the event of a failure.
  - *Recovery level objective (RLO)* is the objective that defines the granularity with which you must be able to recover data — whether you must be able to recover the entire farm, Web application, site collection, site, list or library, or item.

- Availability

  For each component within a farm that is covered by an availability plan, an availability SLA may identify availability as a percentage of uptime, often expressed as the number of nines — that is, the percentage of time that a given system is active and working. For example, a system with a 99.999 uptime percentage is said to have five nines of availability.

  **Note:**
  When calculating availability, most organizations specifically exempt or add hours for planned maintenance activities.

  For more information, see Plan for availability (SharePoint Server 2010).

- Disaster recovery

  For each component within a farm that is covered by a disaster recovery plan, an SLA may identify the recovery point objective and recovery time objective. Different recovery time objectives are often set for different circumstances, for example a local emergency versus a regional emergency.

  For more information, see Plan for disaster recovery (SharePoint Server 2010).

# Related content

| | |
|---|---|
| Resource center | Business Continuity Management for SharePoint Server 2010 (http://go.microsoft.com/fwlink/?LinkId=199235) |
| IT Pro content | Plan for backup and recovery (SharePoint Server 2010)<br>Backup and recovery overview (SharePoint Server |

| | |
|---|---|
| | 2010)<br>Plan to protect content by using recycle bins and versioning (SharePoint Server 2010)<br>Plan for availability (SharePoint Server 2010)<br>Availability configuration (SharePoint Server 2010)<br>Plan for disaster recovery (SharePoint Server 2010) |
| Developer content | Data Protection and Recovery (http://go.microsoft.com/fwlink/?LinkId=199237) |

# Plan to protect content by using recycle bins and versioning (SharePoint Server 2010)

Plan to use recycle bins and versioning in an environment to help users protect and recover their data. Recycle bins and versioning are key components of a business continuity strategy.

**Recycle bins** Users can use recycle bins to retrieve deleted objects. Microsoft SharePoint Server 2010 supports two stages of recycle bins, the first-stage Recycle Bin and the Site Collection — also called the second-stage — Recycle Bin. When Recycle Bins are enabled, users can restore items that are in them, including deleted files, documents, list items, lists, and document libraries.

**Versioning** Users can use versioning to help prevent data loss that is caused by overwriting a document. When a site owner turns on versioning in a document library or a list, the library or list keeps multiple copies of a document, item, or file. In the event of an unwanted change, an overwritten file, or document corruption, the previous version can be easily restored by the user.

In this article:

- Protecting content by using recycle bins
- Protecting content by using versioning

## Protecting content by using recycle bins

SharePoint Server 2010 supports two stages of recycle bins, the first-stage Recycle Bin and the Site Collection, or second-stage, Recycle Bin. The recycle bins are enabled and configured at the Web application level. The recycle bins collect deleted documents and list items. When a list item is deleted, any attachments to the item are also deleted and can be restored from the Recycle Bin.

The Recycle Bins can contain multiple copies of a document that each have the same file name and source. These documents cannot be restored over an existing copy of a document. The Recycle Bins cannot be used to recover previous versions or accidental overwrites of documents — you must use versioning to enable this functionality.

The following table describes how an item is deleted and recovered from the first-stage Recycle Bin and the second-stage Recycle Bin.

| When a user does this | The item is | The item can be restored by |
|---|---|---|
| Deletes an item | Held in the first-stage Recycle Bin until the item is deleted from the Recycle Bin or the item has been in the Recycle Bin longer than the time limit configured for an item to be held in the Recycle | Users or site collection administrators |

| When a user does this | The item is | The item can be restored by |
|---|---|---|
| | Bin. | |
| Deletes an item from the Recycle Bin | Held in the second-stage Recycle Bin | Site collection administrators |

Turning off the Recycle Bin for a Web application empties all Recycle Bins and permanently deletes all items in them.

## First-stage Recycle Bin

The first-stage Recycle Bin is located at the site level and is available to users who have Contribute, Design, or Full Control permissions on a site.When a user deletes an item from a Web site, the item is sent to the site's first-stage Recycle Bin. Items located in the first-stage Recycle Bin count toward the site quota.Items remain in one of the first-stage Recycle Bins in the site until a specified time period has been reached (the default setting is 30 days).

When an item is deleted from the Recycle Bin, the item is sent to the second-stage Recycle Bin.

**Note:**
The time limit for the Recycle Bins applies to the total time after the item was first deleted — not the time spent in either Recycle Bin stage.

## Second stage (Site Collection) Recycle Bin

The second-stage Recycle Bin is located at the site collection administrator level. The second-stage Recycle Bin is organized into two views: objects in the first-stage Recycle Bins of all sites in the site collection, and objects in the second-stage Recycle Bin. When an item is deleted from the first-stage Recycle Bin, it can be recovered only by a site collection administrator from the second-stage Recycle Bin.

Items remain in the second-stage Recycle Bin until a specified time period has been reached (the default setting is 30 days) or until the second-stage Recycle Bin reaches its size limit, at which time the oldest items are deleted. The time limit for the Recycle Bins applies to the total time after the item was initially deleted — not the time spent in either Recycle Bin stage.

When a second-stage Recycle Bin is enabled for a Web application, we recommend that you designate how much disk space is available to the second-stage Recycle Bin as a percentage of the quota allotted to the Web application. Items stored in the second-stage Recycle Bin do not count toward the site quota; however, the size that is specified for the second-stage Recycle Bin increases the total size of the site and the content database that hosts it. If no site quota has been set, there is no limit on the size of the second-stage Recycle Bin.

For example, if you have allotted 100 megabytes (MB) of space for the Web application, allotting a 50 percent quota for the second-stage Recycle Bin allots 50 MB for the second-stage Recycle Bin and

150 MB for the Web application as a whole. You can allot up to 100 percent for the second-stage Recycle Bin quota.

For more information about setting quotas, see

- Plan site maintenance and management (SharePoint Server 2010)
- Create quota templates (SharePoint Server 2010)

For more information about how users can use the Recycle Bin in SharePoint Server 2010, see View, restore, or delete items in the Recycle Bin (http://go.microsoft.com/fwlink/?LinkId=90917&clcid=0x409)

For information about configuring the Recycle Bins, see Configure the Recycle Bin (SharePoint Server 2010).

# Protecting content by using versioning

Versioning addresses the issue of losing data by overwriting a document. It allows the document library to keep multiple copies of the same document. In the event of an unwanted change, an overwrite, or a document corruption, the previous version can easily be restored by the user. Versioning can be enabled at the library or list level. Items and files can be versioned.

Before configuring versioning, be sure to read Plan site maintenance and management (SharePoint Server 2010).

For more information about configuring versioning, see Enable and configure versioning (SharePoint Server 2010).

Administrators must closely manage versioning, because if sites have many versions of files and documents, the sites can become quite large. If you do not restrict the size of sites, your sites can surpass your storage capacity. Farm administrators can manage this issue by establishing service level agreements with site owners and by setting size quotas on sites. For more information about managing versioning, see Manage versioning by using quotas (SharePoint Server 2010).

# Plan for backup and recovery (SharePoint Server 2010)

This article describes the stages involved in planning for backup and recovery, which include determining backup and recovery strategies for a Microsoft SharePoint Server 2010 environment and deciding which tools to use. The stages do not need to be done in the order listed, and the process may be iterative.

When you plan for how you will use backup and recovery for disaster recovery, consider common events, failures, and errors; local emergencies; and regional emergencies.

For detailed information about Microsoft SharePoint Server 2010 backup and recovery, see Backup and recovery overview (SharePoint Server 2010).

In this article:

- Define business requirements
- Choose what to protect and recover in your environment
- Choose tools
- Determine strategies
- Plan for enhanced backup and recovery performance

## Define business requirements

To define business requirements, determine the following for each farm and service in the environment:

- *Recovery point objective (RPO)* is the objective for the maximum amount of time between the last available backup and any potential failure point. It is determined by the amount of data that the business can afford to lose in the event of a failure.
- *Recovery time objective (RTO)* is the objective for the maximum time a data recovery process will take. It is determined by the amount of time the business can afford for the site or service to be unavailable.
- *Recovery level objective (RLO)* is the objective that defines the granularity with which you must be able to recover data — whether you must be able to recover the entire farm, Web application, site collection, site, list or library, or item.

Shorter RPO and RTO, and greater granularity of RLO, all tend to cost more.

A worksheet to help you plan your strategies for backup and recovery for your SharePoint Server 2010 environment can be downloaded from SharePoint 2010 Products backup and recovery planning workbook (http://go.microsoft.com/fwlink/?LinkID=184385).

# Choose what to protect and recover in your environment

Your business requirements will help you determine which components of the environment you need to protect, and the granularity with which you need to be able to recover them.

The following table lists components of a SharePoint environment that you might decide to protect, and the tools that can be used to back up and recover each component.

| Component | SharePoint backup | Microsoft SQL Server 2008 with Service Pack 1 (SP1) and Cumulative Update 2 | System Center Data Protection Manager (DPM) 2010 | File system backup |
|---|---|---|---|---|
| Farm | Yes | | Yes[6] | |
| Service applications | Yes | | | |
| Web application | Yes | | | |
| Content databases | Yes | Yes | Yes | |
| Site collection | Yes[1, 2] | Yes[1, 2] | Yes[1, 2] | |
| Site | Yes[2] | Yes[2] | Yes | |
| Document library or list | Yes[2] | Yes[2] | Yes | |
| List item or document | | | Yes | |
| Content stored in remote BLOB stores | Yes[3] | Yes[3] | Yes[3] | |
| Customizations deployed as solution packages | Yes[7] | Yes[7] | Yes[6, 7] | |
| Changes to Web.config made by using Central Administration or an API | Yes | Yes | Yes[4] | |
| Configuration settings | Yes[2, 8] | Yes[2, 8] | Yes [2, 9] | |

| Component | SharePoint backup | Microsoft SQL Server 2008 with Service Pack 1 (SP1) and Cumulative Update 2 | System Center Data Protection Manager (DPM) 2010 | File system backup |
|---|---|---|---|---|
| (SharePoint) | | | | |
| Customizations not deployed as solution packages | | | Yes. Files can be recovered if protected as files.[4, 5] | Yes |
| Changes to Web.config *not* made by using Central administration or an API | | | Yes[4] | Yes |
| IIS configurations not set through SharePoint | | | Yes[5] | Yes |
| SQL Server Reporting Services databases | | Yes | Yes | |

[1] Farm-level and database-level backup and restore can be used for site collection recovery if a single site collection is stored in a database.

[2] Farm-level and database-level backups can be used with SharePoint Server 2010 unattached database recovery to restore site collections, sites, lists, and configurations.

[3] Content stored in remote BLOB stores is backed up and restored with other content, as long as the Remote BLOB Storage (RBS) provider in use has this capability.

[4] Changes to Web.config can be backed up by using file system backup from DPM 2010.

[5] IIS configurations can be recovered by using a bare metal backup from DPM 2010.

[6] DPM 2010 can recover this item by using a combination of a bare metal backup and SharePoint Server 2010 backup. It cannot be backed up and recovered as an object.

[7] Fully-trusted solution packages are stored in the configuration database, and sandboxed solutions are stored in content databases. They can be recovered as part of farm or content database recovery.

[8] Configuration settings can be recovered from farm-level backups. For more information, see Restore a farm (SharePoint Server 2010).

[9] The Central Administration content database and the configuration database for a SharePoint Server 2010 farm can be recovered but only as part of a full-farm recovery to the same farm, with the same computers.

**Note**

- You can register SharePoint Server 2010 with Windows Server Backup by using the stsadm.exe **-o -registerwsswriter** operation to configure the Volume Shadow Copy Service (VSS) writer for SharePoint Server 2010. Windows Server Backup then includes SharePoint Server 2010 in server-wide backups. When you restore from a Windows Server backup, you can select Microsoft SharePoint Foundation (no matter which version of SharePoint 2010 Products is installed), and all components reported by the VSS writer forSharePoint Server 2010 on that server at the time of the backup will be restored.
- Windows Server Backup is recommended only for use with for single-server deployments.

## Choose what to recover from within SharePoint content databases

From within a content database, you can recover site collections, sites, lists and libraries.

Backup and recovery tools provide different levels of recovery for content within a content database. Recovering an object from within a content database is always more complex than recovering an entire content database.

## Protecting customizations

Customizations to SharePoint sites can include:

- Master pages, page layouts and cascading style sheets. These objects are stored in the content database for a Web application.
- Web Parts, site or list definitions, custom columns, new content types, custom fields, custom actions, coded workflows, or workflow activities and conditions.
- Third-party solutions and their associated binary files and registry keys, such as IFilters.
- Changes to standard XML files.
- Custom site definitions (Webtemp.xml).
- Changes to the Web.config file.

How customizations are deployed, and how changes are made to the Web.config file, have a significant effect on which tools can be used to back up and recover customizations. To provide the greatest opportunity for recovery, we recommend that you deploy customizations by using solution packages and make changes to the Web.config file by using Central Administration or the SharePoint APIs and object model.

## Protecting workflows

Workflows are a special case of customizations that you can back up and recover. Make sure that your backup and recovery plan addresses any of the following scenarios that apply to your environment:

- Declarative workflows, such as those created in Microsoft SharePoint Designer 2010, are stored in the content database for the site collection to which they are they are deployed. Backing up the content database protects these workflows.
- Custom declarative workflow actions have components in the following three locations:

  a. The Visual Studio assemblies for the Activities are stored in the global assembly catalog (GAC).

  b. The XML definition files (.ACTIONS files) are stored in the 14\TEMPLATE\{LCID}\Workflow directory.

  c. An XML entry to mark the activity as an authorized type is stored in the Web.config file for the Web applications in which it is used.

  If your farm workflows use custom actions, you should use a file backup system to protect these files and XML entries. Similar to SharePoint Server 2010 features such as Web parts and event receivers, these files should be reapplied to the farm as needed after recovery.
- Workflows that depend on custom code, such as those that are created by using Visual Studio, are stored in two locations. The Visual Studio assemblies for the workflow are stored in the global assembly catalog (GAC), and the XML definition files are stored in the Features directory. This is the same as other types of SharePoint Server 2010 features such as Web parts and event receivers. If the workflow was installed as part of a solution package, backing up the content database protects these workflows.
- If you create a custom workflow that interacts with a site collection other than the one where the workflow is deployed, you must back up both site collections to protect the workflow. This includes workflows that write to a history list or other custom list in another site collection. Performing a farm backup is sufficient to back up all site collections in the farm and all workflows that are associated with them.
- Workflows that are not yet deployed must be backed up and restored separately like any other data file. When you are developing a new workflow but have not yet deployed it to the SharePoint Server 2010 farm, make sure that you back up the folder where you store your workflow project files by using Windows Backup or another file system backup application.

## Protecting service applications

Service applications in a SharePoint Server 2010 environment can be made up of both service settings and one or more databases, or just service settings. You cannot restore a complete service application by restoring the database only; however, you can restore the databases for a service application and then reprovision the service application. For more information, see Restore a service application (SharePoint Server 2010).

### Protecting SQL Server Reporting Services databases

SharePoint Server 2010 backup and recovery does not include SQL Server Reporting Services databases. You must use SQL Server tools. For more information, see Backup and Restore Operations for a Reporting Services Installation (http://go.microsoft.com/fwlink/?LinkId=186642).

## Choose tools

To choose the right tools for backup and recovery, you need to determine whether you can meet the continuity requirements you have set for your business within your budget for time and resources.

Key factors to consider when choosing tools include:

- Speed of backup: Can the tool perform within the maintenance window for your databases? You should test any backup system to ensure that it meets your needs on your hardware.
- Completeness of recovery.
- Granularity of objects that can be recovered.
- Backup type supported (full, differential, or incremental).
- Complexity of managing the tool.

The following table compares the type of backup and size of farm that can be backed up in a six-hour window for backup and recovery tools available from Microsoft.

| **Tool** | **Backup type** | **Size of backup completed in six hours**[1] |
|---|---|---|
| SharePoint farm backup and recovery | Full, differential | 600 GB |
| SQL Server | Full, differential | 600 GB |
| System Center Data Protection Manager | Incremental | Terabytes |

[1]Backup size was determined by backing up a system that totals the specified size on the test hardware listed in the following section.

**Note:**
The SharePoint Server 2010 and SQL Server backups were performed with backup compression turned on.

### Test hardware

The following table lists the hardware used in the tests that determined the size of backup that could be completed in a six-hour window.

| Component | Description |
|---|---|
| Processor | 64-bit dual processor, 3 GHz |
| RAM | 8 GB |
| Disk | 2 terabyte NTFS file system-formatted partition |
| Network | 100 megabits per second (Mbps) or faster connection between client computers and server |
| Network share | Network share with 1.25 terabytes free space |

**Note:**
The upper size limit for performing SharePoint Server 2010 site collection backups is 85 GB.

For detailed information about the backup and recovery systems that can be used with Microsoft SharePoint Server 2010, see the following resources:

- Backup and recovery overview (SharePoint Server 2010)
- Backing Up and Restoring Databases in SQL Server (http://go.microsoft.com/fwlink/?LinkID=186643)
- Data Protection Manager 2010 Release Candidate Overview (http://go.microsoft.com/fwlink/?LinkID=186655)

# Determine strategies

Based on your business requirements, recovery needs, and the tools you have chosen, determine and document the backup and recovery strategies for your environment.

It is not uncommon for IT departments that support SharePoint Server 2010 environments to decide to use more than one tool to protect the environment, as they determine the strategies that they will use.

For example, in an environment with databases that are managed by DBAs, the strategies in the following list might be employed:

- All databases are backed up by SQL Server. The backup interval that is set for each database is based on the following:
  - The business impact of the content or service.
  - The standard rate of change for the database.
  - The effect on performance that the backup has on the environment.
- Small, rapidly changing, very high-business-impact content databases are additionally protected by SQL Server database snapshots that are stored on a separate physical disk. Only one snapshot is

  stored per database, and snapshots are discarded regularly, so that the effect on performance is minimized. The snapshot interval that is set for each database is based on the following:

  - The business impact of the content or service.
  - The standard rate of change for the database.
  - The effect on performance that the snapshot has on the environment.
  - The amount of space required to store the snapshot.

  Recovering from a snapshot is faster than standard recovery because a snapshot, along with its underlying database, can be treated by SharePoint Server 2010 as an unattached database. However, the process of creating snapshots can decrease the performance of the underlying database. We recommend that the effect that snapshots have on the performance of your system be tested before they are implemented, and that snapshots be discarded regularly to reduce the space required.

  > **Note:**
  > If you are using RBS, and the RBS provider that you are using does not support snapshots, you cannot use snapshots for backup. For example, the SQL FILESTREAM provider does not support snapshots.

- SharePoint Server 2010 backup is used to protect service applications. The backup interval is based on the following:
  - The business impact of the service.
  - The standard rate of change for the database.
  - The effect on performance that the backup has on the database.
- All restore operations are performed through SharePoint Server 2010. The choice of which restore system to use is determined by the type of backup that is available and the object being restored.

Other tools should be part of your business continuity strategy. Consider how you will use Recycle Bins and versioning in site collections throughout the environment. For more information, see Plan for business continuity management (SharePoint Server 2010).

# Plan for enhanced backup and recovery performance

As you plan your backup and recovery strategy, consider the following recommendations to help you decrease the effect of backup and recovery on system performance.

By design, most backup jobs consume as many I/O resources as they can to finish the job in the available time for maintenance; therefore, you might see disk queuing and you might see that all I/O requests come back more slowly than usual. This is typical and should not be considered a problem.

## Follow recommendations for configuring SQL Server and storage

Follow the general recommendations for configuring SQL Server and storage for a SharePoint Server 2010 environment. For more information, see Storage and SQL Server capacity planning and configuration (SharePoint Server 2010).

## Minimize latency between SQL Server and the backup location

In general, it is best to use a local disk, not a network drive, for backups. If you are backing up multiple servers, you may want to have a directly connected computer that both servers can write to. Network drives that have 1 millisecond or less latency between them and the computers that are running SQL Server will perform well. If your farm has multiple servers in it (including the computer that is running SQL Server), you must use UNC network paths for the SharePoint farm backup location.

## Avoid processing conflicts

Do not run backup jobs during times in which users require access to the system.

To avoid I/O bottlenecks, perform the main backup to a separate disk, and only then copy to tape.

Consider staggering backups so that not all databases are backed up at the same time.

SharePoint Server 2010 backups use SQL Server backups. When using compression with your backups, be mindful not to overwhelm SQL Server. For example, some third-party backup tools compress data during backup, which can disrupt SQL Server performance. There are tools available to throttle the compression processes and control the effect on SQL Server.

## Follow SQL Server backup and restore optimization recommendations

If you are running SQL Server 2008 Enterprise, we recommend that you use backup compression. For more information, see Backup Compression (SQL Server) (http://go.microsoft.com/fwlink/?LinkId=179525).

If you are using SQL Server backups, use a combination of full, differential, and transaction log backups for the full recovery model to minimize recovery time. Differential database backups are usually faster to create than full database backups, and they reduce the amount of transaction log required to recover the database.

If you are using the full recovery model in SQL Server 2008, we recommend that you use the truncate option during backup to avoid maintenance issues.

For detailed recommendations about how to optimize SQL Server backup and restore performance, see Optimizing Backup and Restore Performance in SQL Server (http://go.microsoft.com/fwlink/?LinkId=126630).

### Ensure sufficient write performance on the backup drive

Carefully consider whether to use redundant array of independent disks (RAID) on your disk backup device. For example, RAID 5 has low write performance, approximately the same speed as for a single disk. (This is because RAID 5 maintains parity information.) Using RAID 10 for a backup device may provide faster backups. For more information about how to use RAID with backups, see Configure RAID for maximum SQL Server I/O throughput (http://go.microsoft.com/fwlink/?LinkId=126632).

## Related content

| | |
|---|---|
| Resource center | Business Continuity Management for SharePoint Server 2010(http://go.microsoft.com/fwlink/?LinkId=199235) |
| IT Pro content | Backup and recovery overview (SharePoint Server 2010)<br>Backup and recovery (SharePoint Server 2010)<br>Plan for availability (SharePoint Server 2010)<br>Availability configuration (SharePoint Server 2010)<br>Plan for disaster recovery (SharePoint Server 2010) |
| Developer content | Data Protection and Recovery (http://go.microsoft.com/fwlink/?LinkID=199237) |

# Backup and recovery overview (SharePoint Server 2010)

This article describes the backup architecture and recovery processes that are available in Microsoft SharePoint Server 2010, including farm and granular backup and recovery, and recovery from an unattached content database. Backup and recovery operations can be performed through the user interface or through Windows PowerShell cmdlets. Built-in backup and recovery tools may not meet all the needs of your organization.

In this article:

- Backup and recovery scenarios
- Backup architecture
- Recovery processes

## Backup and recovery scenarios

Backing up and recovering data supports many business scenarios, including the following:

- Recovering unintentionally deleted content that is not protected by the Recycle Bin or versioning.
- Moving data between installations as part of a hardware or software upgrade.
- Recovering from an unexpected failure.

## Backup architecture

SharePoint Server 2010 provides two backup systems: farm and granular.

### Farm backup architecture

The farm backup architecture in SharePoint Server 2010 starts a Microsoft SQL Server backup of content and service application databases, writes configuration content to files, and also backs up the Search index files and synchronizes them with the Search database backups.

The following illustration shows the farm backup system.

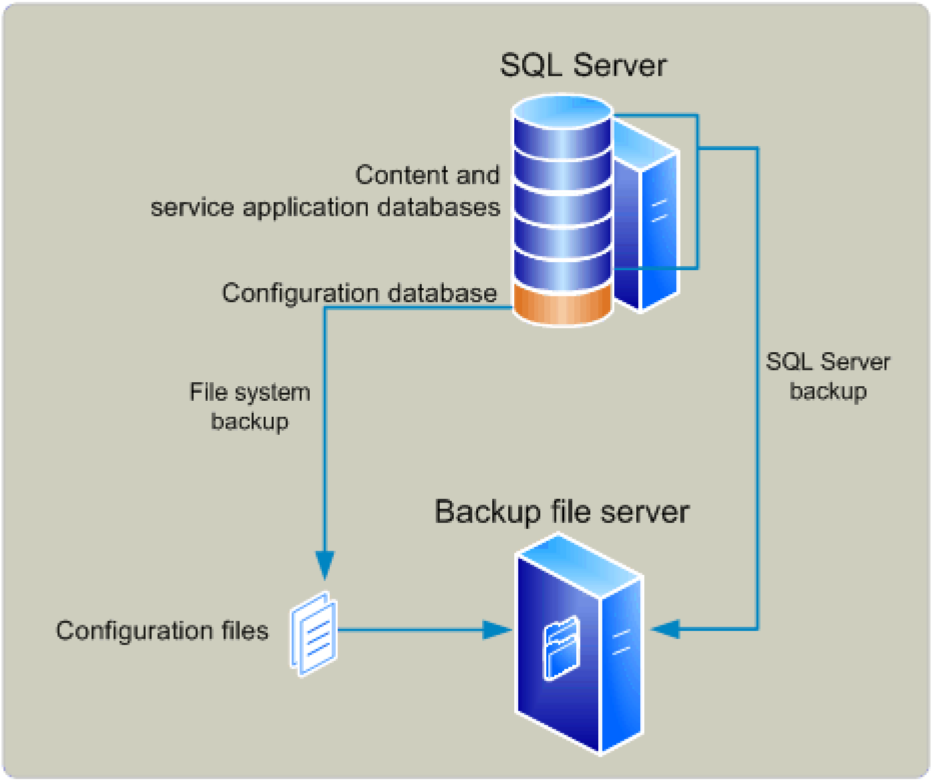

Both full and differential backups are supported. *Full* backups create a new backup of the complete system. *Differential* backups create a backup of all the data that is stored in databases that has changed since the last full backup.

The farm backup system is organized hierarchically. The components in a farm that can be selected for backup include the following:

- **Farm** The farm is the highest-level object. You can select from the following options when you perform a farm backup:
    - Content and configuration data (default)

      The whole server farm is backed up. This includes settings from the configuration database.
    - Configuration-only

      Configuration database settings are backed up so that you can apply configurations across farms. For more information, see Configuration-only backup use and benefits later in this article.
- **Web application** Within a Web application, you can select one or more of the content databases to back up.

  A Web application backup includes the following:
    - Application pool name and application pool account

  - Authentication settings
  - General Web application settings such as alerts and managed paths
  - Internet Information Services (IIS) binding information, such as the protocol type, host header, and port number
  - Changes to the Web.config file that have been made through the object model or Central Administration

    **Note:**
    Changes to the Web.config file that have been made to support claims-based authentication that uses forms-based authentication are not included in backups, because those changes are made manually. For more information, see Considerations for using farm backups later in this article.

  - Sandboxed solutions

  For recommendations about how to protect these settings, see Plan for backup and recovery (SharePoint Server 2010).

- **Services and service applications (not shared)** An example of a service that is not shared is the State Service. Service and service application backups contain the settings for a service or service application and any databases associated with it.

  **Important:**
  Backups of service applications do not include the related proxy. To back up both the service application and the service application proxy, you must either back up the farm or perform two consecutive backups, selecting the service application in one backup, and selecting the associated service application proxy in the second backup.

  Many service application databases cannot be backed up individually from SharePoint Server 2010. To back up service application databases only, you must use SQL Server backup.

- **Proxies for service applications that are not shared**
- **Shared Services** Shared services require both a service application and a service application proxy to run. If you select the Shared Services node, all of the service applications and the related service application proxies on the farm will be backed up.

  **Note:**
  The backup hierarchy enables you to select individual service applications and service application proxies to back up. However, when you select one or all service applications, or one or all proxies, the related objects are not backed up by default. To back up both parts of a specific service, you must either select the Shared Services node or perform two consecutive backups, selecting the service application in one backup, and selecting the associated service application proxy in the second backup.

**Note**

Some settings in the SharePoint Server 2010 environment are not included in a farm backup. They include the following settings that are stored on Web servers:

- Application pool account passwords
- HTTP compression settings
- Time-out settings
- Custom Internet Server Application Programming Interface (ISAPI) filters
- Computer domain membership
- Internet Protocol security (IPsec) settings
- Network Load Balancing settings
- Secure Sockets Layer (SSL) certificates
- Dedicated IP address settings

## Search service application backup process

Backing up and recovering the Search service application is a special case because of the complexity of interactions between the components of the application.

When a backup of the Search service application is started, SharePoint Server 2010 starts a SQL Server backup of the Search administration database, crawl databases, and property databases, and also backs up the index partition files in parallel.

Consider how the backup and recovery processes for the Search service application affect your service-level agreement. For example, consider how pausing all crawls might affect the freshness of search results.

The backup process is as follows:

1. Master merges are paused to preserve the master index.
2. A full database backup starts.
3. The master index is backed up.
4. Crawls are paused. The pause in crawling is much shorter than during a backup of Microsoft Office SharePoint Server 2007 search, and does not last the full duration of the backup process.
5. All shadow indexes are backed up.
6. An incremental database backup starts.
7. Crawls are resumed.
8. Master merges are resumed.

## Configuration-only backup use and benefits

A configuration-only backup extracts and backs up the configuration settings from a configuration database. By using built-in tools, you can back up the configuration of any configuration database, whether it is currently attached to a farm or not. For detailed information about how to back up a configuration, see Back up a farm configuration (SharePoint Server 2010).

A configuration backup can be restored to the same — or any other — server farm. When a configuration is restored, it will overwrite any settings present in the farm that have values that are set in the configuration backup. If any settings present in the farm are not contained in the configuration backup, they will not be changed. For detailed information about how to restore a farm configuration, see Restore a farm configuration (SharePoint Server 2010).

**Note:**
Web application and service application settings are not included in a configuration backup. You can use Windows PowerShell cmdlets to document and copy settings for service applications. For more information, see Document farm configuration settings (SharePoint Server 2010) and Copy configuration settings from one farm to another (SharePoint Server 2010).

Situations in which you might want to restore a configuration from one farm to another farm include the following:

- Replicating a standardized farm configuration to be used throughout an environment.
- Moving configurations from a development or test environment to a production environment.
- Moving configurations from a stand-alone installation to a farm environment.
- Configuring a farm to serve as part of a standby environment.

SharePoint Server 2010 stores the following kinds of settings in the configuration-only backup:

- Antivirus
- Information rights management (IRM)
- Outbound e-mail settings (only restored when you perform an overwrite).
- Customizations deployed as trusted solutions
- Diagnostic logging

## Considerations for using farm backups

Consider the following before you use farm backups:

- There is no built-in scheduling system for backups. To schedule a backup, we recommend that you create a backup script by using Windows PowerShell, and then use Windows Task Scheduler to run the backup script on a regular basis.
- We do not recommend that you use IIS metabase backup to protect IIS settings. Instead, document all IIS configurations for each Web server by using a tool that provides the configuration monitoring you want, such asMicrosoft System Center Configuration Manager 2010.
- SharePoint Server 2010 backup and recovery can be run together with SQL Server Enterprise features such as backup compression and transparent data encryption.

If you are running SQL Server Enterprise, we strongly recommend that you use backup compression. For more information about backup compression, see Backup Compression (SQL Server) http://go.microsoft.com/fwlink/?LinkID=129381).

If you decide to run databases with transparent data encryption, you must manually back up the key and restore the key — SharePoint Server 2010 backup and restore will not remind you about the key. For more information about transparent data encryption, see Understanding Transparent Data Encryption (TDE) (http://go.microsoft.com/fwlink/?LinkID=129384).

- If a content database is set to use the SQL FILESTREAM remote BLOB storage (RBS) provider, the RBS provider must be installed both on the database server that is being backed up and on the database server that is being recovered to.
- SharePoint Server 2010 backup does not protect:
  - Changes to the Web.config file on Web servers that are not made through Central Administration or the object model.
  - Customizations to a site that are not deployed as part of a trusted or sandboxed solution.
- If you are sharing service applications across farms, be aware that trust certificates that have been exchanged are not included in farm backups. You must back up the certificate store separately or keep the certificates in a separate location. When you restore a farm that shares a service application, you must import and redeploy the certificates and then re-establish any inter-farm trusts.

  For more information, see Exchange trust certificates between farms (SharePoint Server 2010).
- When you restore a farm or Web application that is configured to use any kind of claims-based authentication, duplicate or additional providers may appear to be enabled. If duplicates appear, you must manually save each Web application zone to remove them.
- Additional steps are required when you restore a farm that contains a Web application that is configured to use forms-based authentication. You must re-register the membership and role providers in the Web.config file, and then redeploy the providers. You must perform these steps whether you are restoring at the Web application level or at the farm level.

  For more information, see Back up a Web application (SharePoint Server 2010), Plan authentication methods (SharePoint Server 2010) and Configure claims authentication (SharePoint Server 2010).

## Granular backup and export architecture

The granular backup and export architecture uses Transact-SQL queries and export calls. Granular backup and export is a more read-intensive and processing-intensive operation than farm backup.

From the granular backup system, a user can back up a site collection, or export a site or list.

**Note:**
Workflows are not included in exports of sites or lists.

If you are running SQL Server Enterprise, the granular backup system can optionally use SQL Server database snapshots to ensure that data remains consistent while the backup or export is in progress. When a snapshot is requested, a SQL Server database snapshot of the appropriate content database is taken, SharePoint Server 2010 uses it to create the backup or export package, and then the snapshot is deleted. Database snapshots are linked to the source database where they originated. If the source database goes offline for any reason, the snapshot will be unavailable. For more information about database snapshots, see Database Snapshots (http://go.microsoft.com/fwlink/?LinkId=166158).

Benefits of backing up a site collection by using a snapshot include the following:

- The snapshot ensures that the data that is being read remains consistent while the operation is being performed.
- Users can continue to interact with the site collection while it is being backed up from the database snapshot. This includes adding, editing, and deleting content. However, the changes that users make to the live site will not be included in the site collection backup because the backup is based on the database snapshot.

However, database snapshots can adversely affect performance. For more information about database snapshots and performance, see Limitations and Requirements of Database Snapshots (http://go.microsoft.com/fwlink/?LinkId=166159).

You can use granular backup and export for content that is stored in a database that is configured to use the SQL FILESTREAM RBS provider.

**Note:**
If the RBS provider that you are using does not support snapshots, you cannot use snapshots for content deployment or backup. For example, the SQL FILESTREAM provider does not support snapshots.

**Note:**
We do not recommend that you use SharePoint Server 2010 site collection backup for site collections larger than 85 GB.

The following illustration shows the granular backup and export system.

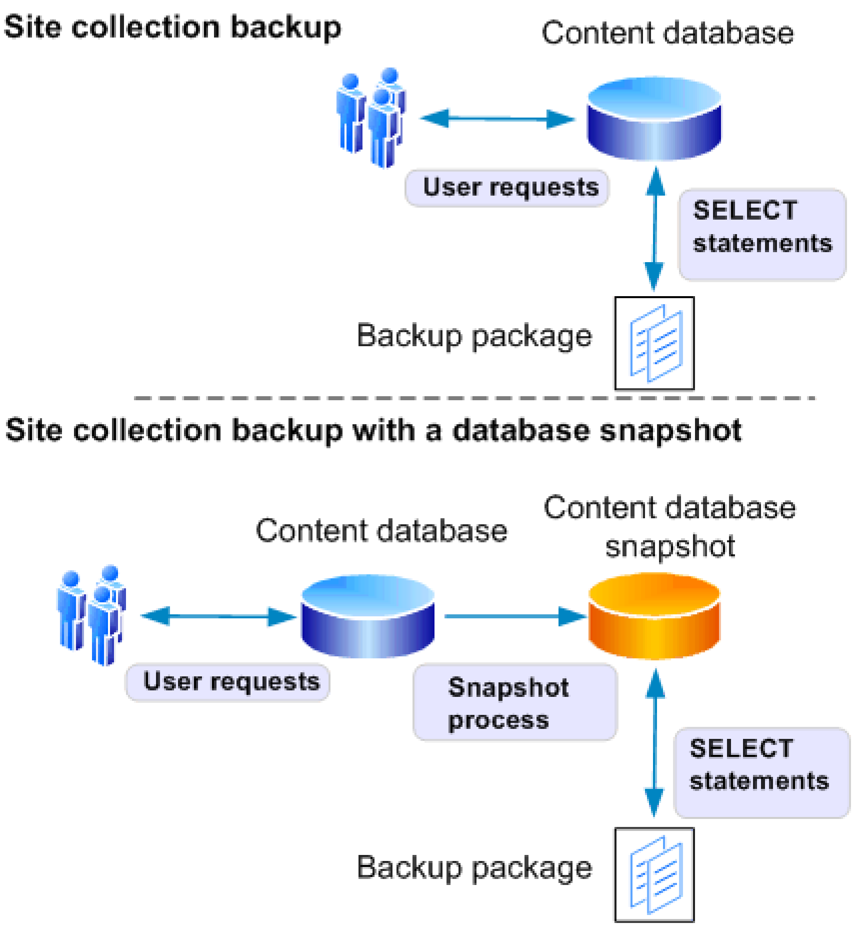

# Recovery processes

SharePoint Server 2010 supports the following primary, built-in recovery options:

- Restore from a farm backup that was created by using built-in tools, or restore from the backup of a component taken by using the farm backup system.
- Restore from a site collection backup.
- Connect to a content database by using the unattached content database feature, back up or export data from it, and then restore or import the data.

## Restoring from a farm backup

Items that can be recovered from a farm backup include the following:

- Farm
    - Content and configuration data (default)

    The whole server farm is restored. This includes settings from the configuration database, and trusted solution packages.

  - Configuration-only

    Only the configuration data is restored. This overwrites any settings in the farm that have values that are set within the configuration-only backup.

- Web applications

  Restores Web applications.

- Service applications

  Restores service applications. Service application recovery can be complex because SharePoint Server 2010 cannot fully reconfigure service application proxies during the restore process. Service application proxies are restored, but are not put in proxy groups. Therefore, they are not associated with any Web applications. For more information about how to restore a Search service application, see Search service application recovery process. For specific information about the operations involved in restoring specific service applications, see Restore a service application (SharePoint Server 2010).

- Content databases

  When content databases are restored, the sandboxed solutions associated with the related site collections are also restored.

## Restoring as new versus restoring as overwrite

By default, SharePoint Server 2010 recovery restores any object as a new instance of the object, instead of overwriting any existing instances with the same name.

When you restore a farm or object as new, the following objects will not work without adjustments, because all GUIDs for objects are assigned new values:

- **Farm.** When you restore a farm as new, you must do the following:
  - Re-create alternate access mapping settings. SharePoint Server 2010 recovery only restores the Default zone of the Web application.
  - Reconfigure settings for any Business Connectivity Services and Managed Metadata service application external sources.
  - Re-associate service application proxies with proxy groups because service application proxies are not assigned to proxy groups when restored. All Web applications will be associated with the default proxy group. You must associate Web applications with other proxy groups if you want to do that.
- Web application.
  - If the Web application name and URL that you provide match a Web application name and URL that already exist in the farm, SharePoint Server 2010 recovery combines them.
  - If you do not want to combine Web applications, you must rename the Web application when you restore it as new.

  - When you restore a Web application as new in the same environment but do not combine Web applications, many other parameters and objects must also be changed. For example, you may have to provide different database file paths and different database names.
- Service applications and service application proxies
  - If you recover a service application and also recover the related service application proxy, you must associate the service application proxy with a proxy group.
  - If you recover a service application and do not also recover the related service application proxy, you must re-create the service application proxy.

  **Note:**
  You cannot restore a service application as new in the same farm. You can restore a service application as new in another farm.

When you restore an object and overwrite the existing object, no changes are necessary.

### Search service application recovery process

The recovery process for the Search service application varies depending on whether you are restoring as new or restoring as overwrite. When you restore as overwrite, no additional steps are necessary.

The restore as new process is as follows:

1. Restore the service application as new, and specify the new farm topology information as you restore.
2. Restore the service application proxy as new. If you did not restore the service application proxy, you must create a new service application proxy and associate it with the Search service application.
3. Associate the service application proxy with the appropriate proxy group and associate the proxy group (if it is not the default proxy group) with the appropriate Web application.
4. For least-privilege deployments, start the Search service and the Search admin query Web service with the appropriate account.

For more information about how to recover the Search service application, see Restore search (SharePoint Server 2010).

## Restoring from a site collection backup

Only site collections can be recovered from a site collection backup.

## Recovering from an unattached content database

SharePoint Server 2010 provides the ability to connect to, and back up from, a content database that is attached to an instance of SQL Server but is not associated with a local SharePoint Web application. Unattached databases that you can connect to include read-only content databases that have been restored from any supported backup technology and SQL Server database snapshots of content databases.

Recovery is the following two-stage process:

1. Back up or export the object from the unattached content database.
2. Restore or import the output of the prior step into SharePoint Server 2010.

The following items can be backed up or exported from an unattached database by using granular backup and export, and then restored:

- Site collection

  Back up by using site collection backup, and then recover by using a site collection restore.
- Site

  Export, and then import.
- Lists and libraries

  Export, and then import.

You can use import to recover content that you backed up from a database configured to use the SQL FILESTREAM RBS provider. The recovered content will be stored by SharePoint Server 2010 using the currently defined storage provider for that content database — that is, if the content database is not set to use RBS, the data will be stored in the content database; if the content database is set to use RBS, the data will be stored in RBS.

# Related content

| | |
|---|---|
| Resource center | Business Continuity Management for SharePoint Server 2010 (http://go.microsoft.com/fwlink/?LinkID=199235) |
| IT pro content | Plan for backup and recovery (SharePoint Server 2010)<br>Backup and recovery (SharePoint Server 2010) |
| Developer content | Data Protection and Recovery (http://go.microsoft.com/fwlink/?LinkID=199237) |

# Plan for availability (SharePoint Server 2010)

This article describes key decisions in choosing availability strategies for a Microsoft SharePoint Server 2010 environment.

As you carefully review your availability requirements, be aware that the higher the level of availability and the more systems that you protect, the more complex and costly your availability solution is likely to be.

Not all solutions in an organization are likely to require the same level of availability. You can offer different levels of availability for different sites, different services, or different farms.

In this article:

- Availability overview
- Choosing an availability strategy and level
- Redundancy and failover between closely located data centers configured as a single farm ("stretched" farm)

## Availability overview

Availability is the degree to which a SharePoint Server 2010 environment is perceived by users to be available. An available system is a system that is resilient — that is, incidents that affect service occur infrequently, and timely and effective action is taken when they do occur.

Availability is part of business continuity management (BCM), and is related to backup and recovery and disaster recovery. For more information about these related processes, see Plan for backup and recovery (SharePoint Server 2010) and Plan for disaster recovery (SharePoint Server 2010).

**Note:**
When calculating availability, most organizations specifically exempt or add hours for planned maintenance activities.

One of the most common measures of availability is percentage of uptime expressed as *number of nines* — that is, the percentage of time that a given system is active and working. For example, a system with a 99.999 uptime percentage is said to have five nines of availability.

The following table correlates uptime percentage with calendar time equivalents.

| Acceptable uptime percentage | Downtime per day | Downtime per month | Downtime per year |
|---|---|---|---|
| 95 | 72.00 minutes | 36 hours | 18.26 days |
| 99 (two nines) | 14.40 minutes | 7 hours | 3.65 days |

| Acceptable uptime percentage | Downtime per day | Downtime per month | Downtime per year |
|---|---|---|---|
| 99.9 (three nines) | 86.40 seconds | 43 minutes | 8.77 hours |
| 99.99 (four nines) | 8.64 seconds | 4 minutes | 52.60 minutes |
| 99.999 (five nines) | 0.86 seconds | 26 seconds | 5.26 minutes |

If you can make an educated guess about the number of total hours downtime you are likely to have per year, you can use the following formulas to calculate the uptime percentage for a year, a month, or a week:

## Costs of availability

Availability is one of the more expensive requirements for a system. The higher the level of availability and the more systems that you protect, the more complex and costly an availability solution is likely to be. When you invest in availability, costs include the following:

- Additional hardware and software, which can increase the complexity of interactions among software applications and settings.
- Additional operational complexity.

The costs of improving availability should be evaluated in conjunction with your business needs — not all solutions in an organization are likely to require the same level of availability. You can offer different levels of availability for different sites, different services, or different farms.

Availability is a key area in which information technology (IT) groups offer service level agreements (SLAs) to set expectations with customer groups. Many IT organizations offer various SLAs that are associated with different chargeback levels.

## Determining availability requirements

To gauge your organization's tolerance of downtime for a site, service, or farm, answer the following questions:

- If the site, service, or farm becomes unavailable, will employees be unable to perform their expected job responsibilities?
- If the site, service, or farm becomes unavailable, will business and customer transactions be stopped, leading to loss of business and customers?

If you answered yes to either of these questions, you should invest in an availability solution.

# Choosing an availability strategy and level

You can choose among many approaches to improve availability in a SharePoint Server 2010 environment, including the following:

- Improve the fault tolerance of server hardware components.
- Increase the redundancy of server roles within a farm.

## Hardware component fault tolerance

Hardware component fault tolerance is the redundancy of hardware components and infrastructure systems such as power supplies at the server level. When planning for hardware component fault tolerance, consider the following:

- Complete redundancy of every component within a server may be impossible or impractical. Use additional servers for additional redundancy.
- Ensure that servers have multiple power supplies connected to different power sources for maximum redundancy.

In any system, we recommend that you work with hardware vendors to obtain fault-tolerant hardware that is appropriate for the system, including redundant array of independent disks (RAID) arrays.

## Redundancy within a farm

SharePoint Server 2010 supports running server roles on redundant computers (that is, scaling out) within a farm to increase capacity and to provide basic availability.

The capacity that you require determines both the number of servers and the size of the servers in a farm. After you have met your base capacity requirements, you may want to add more servers to increase overall availability. The following illustration shows how you can provide redundancy for each server role.

**Availability within a server farm**

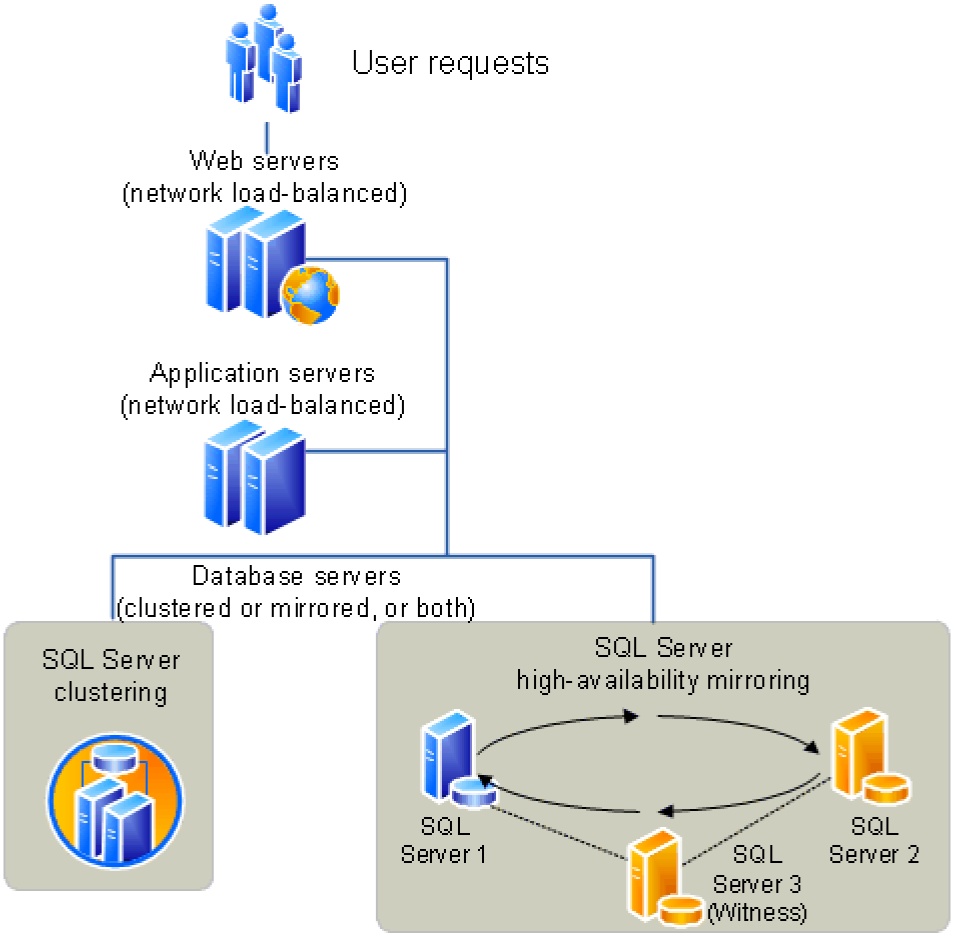

The following table describes the server roles in a SharePoint Server 2010 environment and the redundancy strategies that can be used for each within a farm.

| **Server role** | **Preferred redundancy strategy within a farm** |
|---|---|
| Front-end Web server | Deploy multiple front-end Web servers within a farm, and use Network Load Balancing (NLB). |
| Application server | Deploy multiple application servers within a farm. |
| Database server | Deploy database servers by using clustering or high-availability database mirroring. |

## Database availability strategies

You can use Microsoft SQL Server failover clustering or SQL Server high-availability database mirroring to support availability of databases in a SharePoint Server 2010 environment.

### SQL Server failover clustering

Failover clustering can provide availability support for an instance of SQL Server. A failover cluster is a combination of one or more nodes or servers, and two or more shared disks. A failover cluster instance appears as a single computer, but has functionality that provides failover from one node to another if the current node becomes unavailable. SharePoint Server 2010 can run on any combination of active and passive nodes in a cluster that is supported by SQL Server.

SharePoint Server 2010 references the cluster as a whole; therefore, failover is automatic and seamless from the perspective of SharePoint Server 2010.

For detailed information about failover clustering, see Getting Started with SQL Server 2008 Failover Clustering (http://go.microsoft.com/fwlink/?LinkID=102837&clcid=0x409) and Configure availability by using SQL Server clustering (SharePoint Server 2010).

### SQL Server high-availability mirroring

Database mirroring is a SQL Server technology that can deliver database redundancy on a per-database basis. In database mirroring, transactions are sent directly from a principal database and server to a mirror database and server when the transaction log buffer of the principal database is written to disk. This technique can keep the mirror database almost up to date with the principal database. SQL Server Enterprise Edition provides additional functionality that improves database mirroring performance. For more information, see SQL Server 2008 R2 and SharePoint 2010 Products: Better Together (white paper) (SharePoint Server 2010).

For mirroring within a SharePoint Server 2010 farm, you must use high-availability mirroring, also known as high-safety mode with automatic failover. High-availability database mirroring involves three server instances: a principal, a mirror, and a witness. The witness server enables SQL Server to automatically fail over from the principal server to the mirror server. Failover from the principal database to the mirror database typically takes several seconds.

A change from previous versions is that SharePoint Server 2010 is mirroring-aware. After you have configured a database mirror instance of SQL Server, you then use SharePoint Central Administration or Windows PowerShell cmdlets to identify the failover (mirror) database server location for a configuration database, content database, or service application database. Setting a failover database location adds a parameter to the connection string that SharePoint Server 2010 uses to connect to SQL Server. In the event of a SQL Server time-out event, the following occurs:

1. The witness server that is configured for SQL Server mirroring automatically swaps the roles of the primary and mirror databases.
2. SharePoint Server 2010 automatically attempts to contact the server that is specified as the failover database.

For information about how to configure database mirroring, see Configure availability by using SQL Server database mirroring (SharePoint Server 2010).

For general information about database mirroring, see Database Mirroring (http://go.microsoft.com/fwlink/?LinkID=180597).

**Note:**
Databases that have been configured to use the SQL Server FILESTREAM remote BLOB store provider cannot be mirrored.

### Comparison of database availability strategies for a single farm: SQL Server failover clustering vs. SQL Server high-availability mirroring

The following table compares failover clustering to synchronous SQL Server high-availability mirroring.

| | **SQL Server failover clustering** | **SQL Server high-availability mirroring** |
|---|---|---|
| Time to failover | Cluster member takes over immediately upon failure. | Mirror takes over immediately upon failure. |
| Transactional consistency? | Yes | Yes |
| Transactional concurrency? | Yes | Yes |
| Time to recovery | Shorter time to recovery (milliseconds) | Slightly longer time to recovery (milliseconds). |
| Steps required for failover? | Failure is automatically detected by database nodes; SharePoint Server 2010 references the cluster so that failover is seamless and automatic. | Failure is automatically detected by the database; SharePoint Server 2010 is aware of the mirror location, if it has been configured correctly, so that failover is automatic. |
| Protection against failed storage? | Does not protect against failed storage, because storage is shared between nodes in the cluster. | Protects against failed storage because both the principal and mirror database servers write to local disks. |
| Storage types supported | Shared storage (more expensive). | Can use less-expensive direct-attached storage (DAS). |
| Location requirements | Members of the cluster must be on the same subnet. | Principal, mirror, and witness servers must be on the same LAN (up to 1 millisecond latency roundtrip). |
| Recovery model | SQL Server full recovery model recommended. You can use the | Requires SQL Server full recovery model. |

| | SQL Server failover clustering | SQL Server high-availability mirroring |
|---|---|---|
| | SQL Server simple recovery model, but the only available recovery point if the cluster is lost will be the last full backup. | |
| Performance overhead | Some decrease in performance may occur while a failover is occurring. | High-availability mirroring introduces transactional latency because it is synchronous. It also requires additional memory and processor overhead. |
| Operational burden | Set up and maintained at the server level. | The operational burden is larger than clustering. Must be set up and maintained for all databases. Reconfiguring after failover is manual. |

## Service application redundancy strategies

The redundancy strategy you follow for protecting service applications that run in a farm varies, depending on where the service application stores data.

### Service applications that store data outside a database

To protect service applications that store data outside a database, install the service application on multiple application servers to provide redundancy within the environment.

In this release of SharePoint Server 2010, when you install a service application on multiple application servers, the timer jobs run either on all the application servers that are running the service instance associated with that service application or on the first available server. If an application server fails, timer jobs that are running on that server will be restarted on another server when the next timer job is scheduled to run.

Installing a service application on multiple application servers keeps the service application running, but does not guarantee against data loss. If an application server fails, the active connections for that application server will be lost and users will lose some data.

The following service applications store data outside a database:

- Access Services
- Excel Services Application

### Service applications that store data in databases

To help protect service applications that store data in databases, you must follow these steps:

1. Install the service on multiple application servers to provide redundancy within the environment.
2. Configure SQL Server clustering or mirroring to protect the data.

The following service applications store data in databases:

- Search service application, including the following databases:
    - Search Administration
    - Crawl
    - Property

        **Note:**
        Mirroring the Search databases is supported, but providing redundancy for Search requires additional work. For details, see the section Search redundancy strategies within a farm.

- User Profile service, including the following databases:
    - Profiles
    - Social
    - Synchronization

        **Note:**
        Mirroring the Synchronization database is not supported.

- Business Data Connectivity service application
- Application Registry service application

    We do not recommend mirroring the Application Registry database, because it is only used when upgrading Microsoft Office SharePoint Server 2007 Business Data Catolog information to SharePoint Server 2010.

- Usage and Health Data Collection service application

    **Note:**
    We recommend that you do not mirror the Usage and Health Data Collection service application Logging database.

- Managed Metadata service application
- Secure Store service application
- State service application
- Web Analytics service application, including the following databases:
    - Reporting
    - Staging

        **Note:**

Mirroring the Staging database is not supported.

- Word Automation Services service application
- Microsoft SharePoint Foundation Subscription Settings Service
- PerformancePoint Services

## Search redundancy strategies within a farm

### Server Only

The Search service application is a special case for redundancy within a farm. The following illustration shows how redundancy and failover can be configured for a medium dedicated Search service application that crawls approximately 40 million items. For more information about the architecture of the Search service application, see "Search Architectures for Microsoft SharePoint Server 2010" in the article Technical diagrams (SharePoint Server 2010).

**Redundant Search service application**

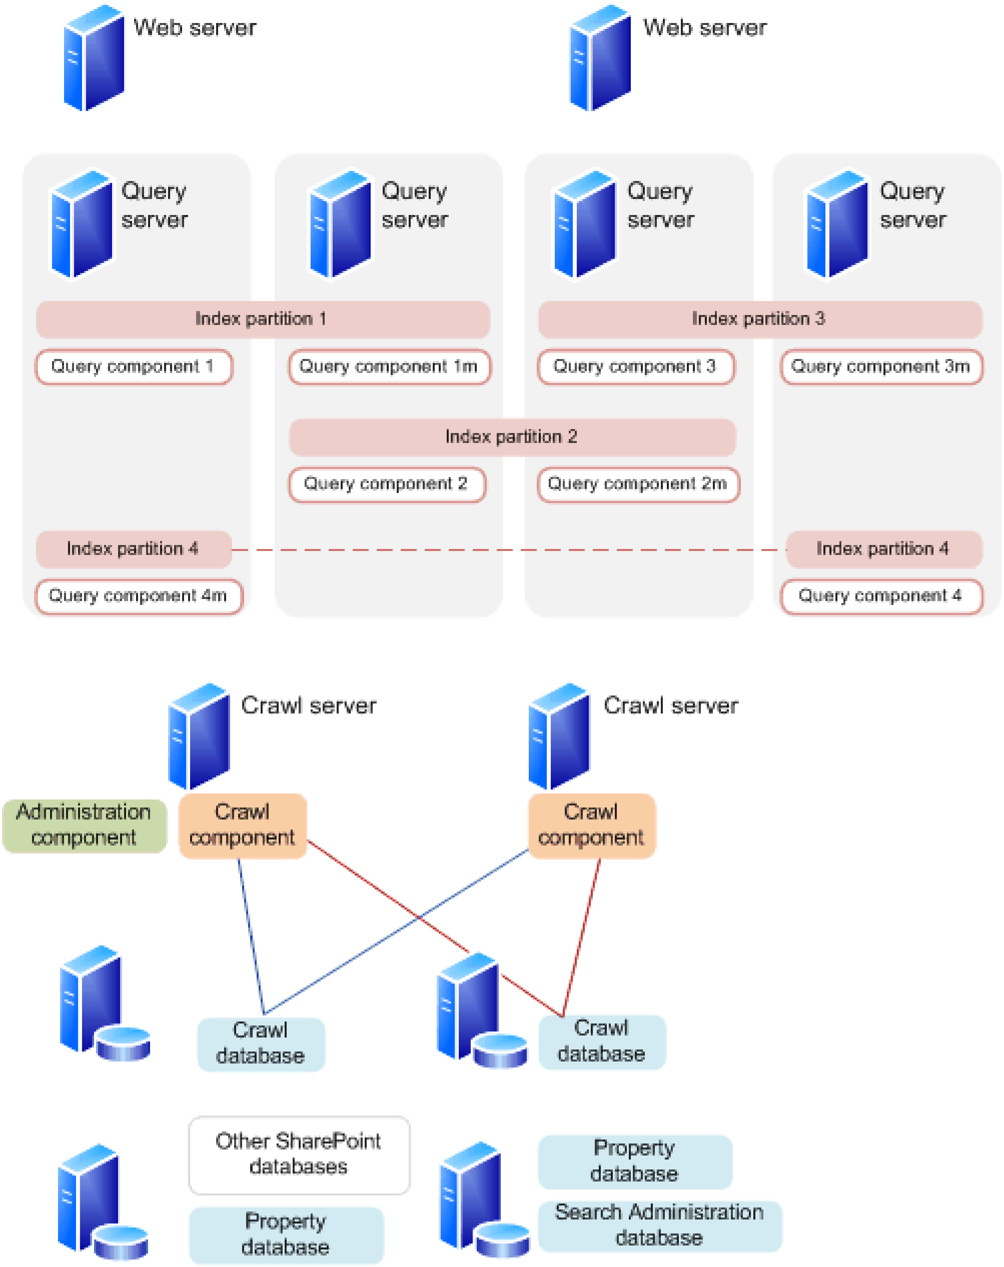

- Query server. A query server hosts query components and index partitions.
  - *Query components* return search results. Each query component is part of an index partition, which is associated with a specific property database that contains metadata associated with a

    specific set of crawled content. You can make an index partition redundant by adding "mirror" query components to an index partition and putting them on different farm servers.

    **Note:**
    The use of the term *mirror query components* refers to identical file copies, not to SQL Server database mirroring.

  - *Index partitions* are groups of query components, each of which holds a subset of the full text index and returns search results. Each index partition is associated with a specific property database that contains metadata that is associated with a specific set of crawled content. You can decide which servers in a farm will handle queries by creating a query component on that server. If you want to balance the load of handling queries across multiple farm servers, add query components to an index partition and associate them with the servers that you want to use to handle queries. For more information, see Add or remove a query component. You can make an index partition redundant by adding mirror query components to an index partition and putting them on different query servers.
- Crawl server. A crawl server hosts crawl components and a search administration component.
  - *Crawl components* process crawls of content sources, propagate the resulting index files to query components, and add information about the location and crawl schedule of content sources to their associated crawl databases. Crawl components are associated with a single Search service application. You can distribute the crawl load by adding crawl components to different crawl servers. You can have as many crawl components on a given crawl server as resources allow. If you have many content locations, you can add crawl components and crawl databases and dedicate them to specific content. Each crawl component on a given crawl server should be associated with a separate crawl database. For redundancy, we recommend that you have at least two crawl components. Each crawl component should be set to crawl both crawl databases. If a database grows to more than 25 million items, we recommend that you add a new crawl database and crawl component.
  - The *search administration component* monitors incoming user actions and updates the search administration database. Only one search administration component is allowed per Search service application.The search administration component can run on any server, preferably either a crawl server or a query server.
- Database servers. Database servers host crawl databases, property databases, the search administration database, and other SharePoint Server 2010 databases.
  - Crawl database

    Crawl databases contain data that is related to the location of content sources, crawl schedules, and other information that is specific to crawl operations for a specific Search service application. You can distribute the database load by adding crawl databases to different computers that are running SQL Server. Crawl databases are associated with crawl components and can be dedicated to specific hosts by creating host distribution rules. For more information about crawl components, see Add or remove a crawl component. For more information about host distribution rules, see Add or remove a host distribution rule. Crawl databases are redundant if they are mirrored or deployed to a SQL Server failover cluster.

- Property database

  Property databases contain metadata that is associated with crawled content. You can distribute the database load of queries by adding property databases to different computers that are running SQL Server. Property databases are associated with index partitions and return any metadata associated with content in query results.

  Property databases are redundant if they are mirrored or deployed to a SQL Server failover cluster.

- Search Administration database

  There is only one Search Administration database per Search service application instance in a farm.

  The Search Administration database is only redundant if it is mirrored or deployed to a SQL Server failover cluster.

For more information about search redundancy, see Manage search topology.

# Redundancy and failover between closely located data centers configured as a single farm ("stretched" farm)

Some enterprises have data centers that are located close to one another with high-bandwidth connections so that they can be configured as a single farm. This is called a *"stretched" farm*. For a stretched farm to work, there must be less than 1 millisecond latency between SQL Server and the front-end Web servers in one direction, and at least 1 gigabit per second bandwidth.

In this scenario, you can provide fault tolerance by following the standard guidance for making databases and service applications redundant.

The following illustration shows a stretched farm.

**Stretched farm**

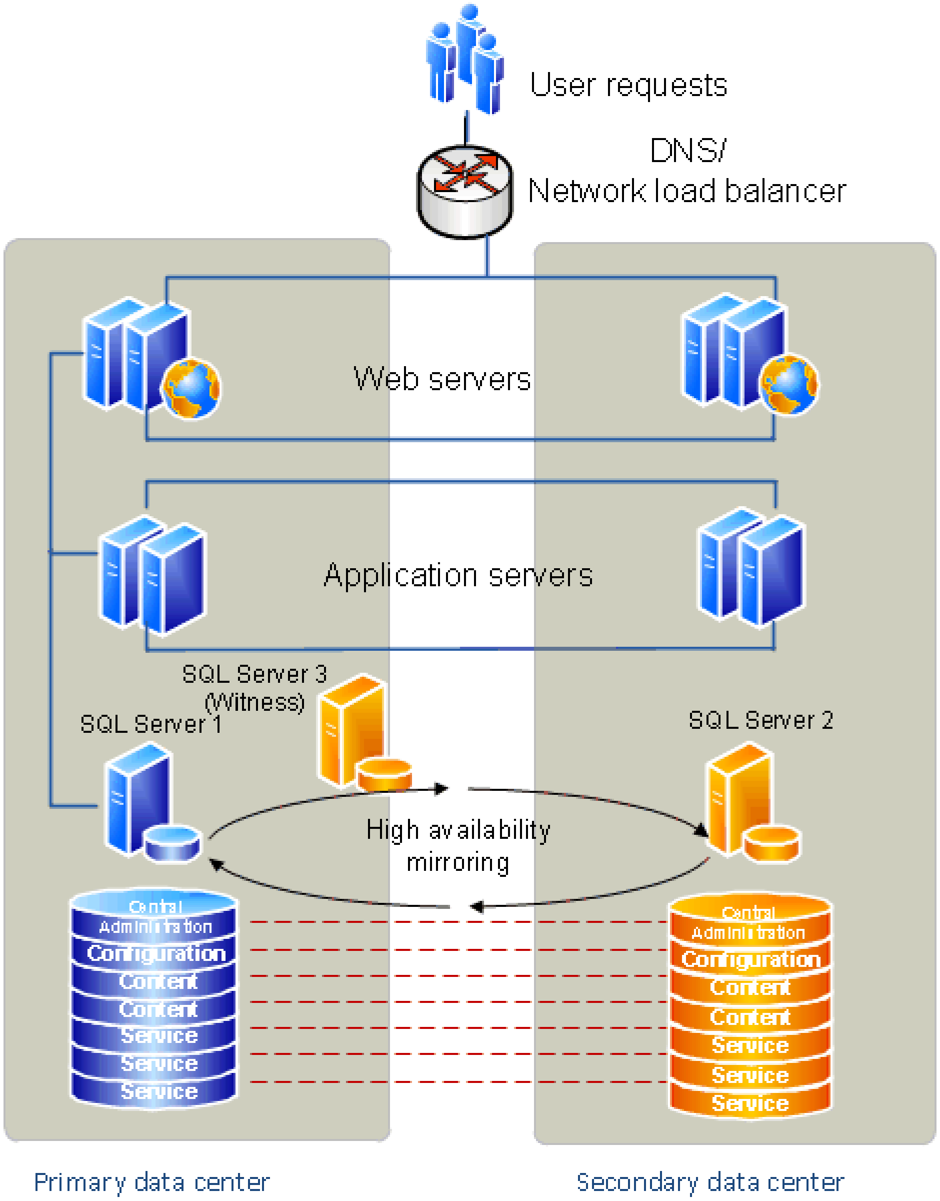

# Plan for disaster recovery (SharePoint Server 2010)

This article describes key decisions in choosing disaster recovery strategies for a Microsoft SharePoint Server 2010 environment.

In this article:

- Disaster recovery overview
- Choose a disaster recovery strategy
- Planning for cold standby data centers
- Planning for warm standby data centers
- Planning for hot standby data centers
- System requirements for disaster recovery

## Disaster recovery overview

For the purposes of this article, we define disaster recovery as the ability to recover from a situation in which a data center that hosts SharePoint Server 2010 becomes unavailable.

The disaster recovery strategy that you use for SharePoint Server 2010 must be coordinated with the disaster recovery strategy for the related infrastructure, including Active Directory domains, Exchange Server, and Microsoft SQL Server. Work with the administrators of the infrastructure that you rely on to design a coordinated disaster recovery strategy and plan.

The time and immediate effort to get another farm up and running in a different location is often referred to as a hot, warm, or cold standby. Our definitions for these terms are as follows:

**Hot standby** A second data center that can provide availability within seconds or minutes.

**Warm standby** A second data center that can provide availability within minutes or hours.

**Cold standby** A second data center that can provide availability within hours or days.

Disaster recovery can be one of the more expensive requirements for a system. The shorter the interval between failure and availability and the more systems you protect, the more complex and costly a disaster recovery solution is likely to be. When you invest in hot or warm standby data centers, costs include:

- Additional hardware and software, which often increase the complexity of operations between software applications, such as custom scripts for failover and recovery.
- Additional operational complexity.

The costs of maintaining hot or warm standby data centers should be evaluated based on your business needs. Not all solutions within an organization are likely to require the same level of availability after a disaster. You can offer different levels of disaster recovery for different content, services, or farms — for example, content that has high impact on your business, or search services, or an Internet publishing farm.

Disaster recovery is a key area in which information technology (IT) groups offer service level agreements (SLAs) to set expectations with customer groups. Many IT organizations offer a variety of SLAs that are associated with different chargeback levels.

When you implement failover between server farms, we recommend that you first deploy and tune the core solution within a farm, and then implement and test disaster recovery.

# Choose a disaster recovery strategy

You can choose among many approaches to provide disaster recovery for a SharePoint Server 2010 environment, depending on your business needs. The following examples show why companies might choose cold, warm, or hot standby disaster recovery strategies.

- Cold standby disaster recovery strategy: A business ships backups to support bare metal recovery to local and regional offsite storage on a regular basis, and has contracts in place for emergency server rentals in another region.

  Pros:

  - Often the cheapest option to maintain, operationally.
  - Often an expensive option to recover, because it requires that physical servers be configured correctly after a disaster has occurred.

  Cons: The slowest option to recover.

- Warm standby disaster recovery strategy: A business ships virtual server images to local and regional disaster recovery farms.

  Pros: Often relatively inexpensive to recover, because a virtual server farm can require little configuration upon recovery.

  Cons: Can be very expensive and time consuming to maintain.

- Hot standby disaster recovery strategy: A business runs multiple data centers, but serves content and services through only one data center.

  Pros: Often relatively fast to recover.

  Cons: Can be quite expensive to configure and maintain.

**Important:**
No matter which disaster recovery solution you decide to implement for your environment, you are likely to incur some data loss.

# Planning for cold standby data centers

In a cold standby disaster recovery scenario, you can recover by setting up a new farm in a new location, (preferably by using a scripted deployment), and restoring backups. Or, you can recover by restoring a farm from a backup solution such as Microsoft System Center Data Protection Manager

2007 that protects your data at the computer level and lets you restore each server individually. This article does not contain detailed instructions for how to create and recover in cold standby scenarios. For more information, see:

- Restore a farm (SharePoint Server 2010)
- Restore customizations (SharePoint Server 2010)

# Planning for warm standby data centers

In a warm standby disaster recovery scenario, you can create a warm standby solution by making sure that you consistently and frequently create virtual images of the servers in your farm that you ship to a secondary location. At the secondary location, you must have an environment available in which you can easily configure and connect the images to re-create your farm environment.

This article does not contain detailed instructions for creating warm standby solutions. For more information about how to plan to deploy farms by using virtual solutions, see Plan for virtualization (SharePoint Server 2010).

# Planning for hot standby data centers

In a hot standby disaster recovery scenario, you can set up a failover farm to provide disaster recovery in a separate data center from the primary farm. An environment that has a separate failover farm has the following characteristics:

- A separate configuration database and Central Administration content database must be maintained on the failover farm.
- All customizations must be deployed on both farms.

  **Note:**
  We recommend that you use scripted deployment to create the primary and failover farm by using the same configuration settings and customizations. For more information, see Install SharePoint Server 2010 by using Windows PowerShell.

- Updates must be applied to both farms, individually.
- SharePoint Server 2010 content databases can be successfully asynchronously mirrored or log-shipped to the failover farm.

  **Note:**
  SQL Server mirroring can only be used to copy databases to a single mirror server, but you can log-ship to multiple secondary servers.

- Service applications vary in whether they can be log-shipped to a farm. For more information, see Service application redundancy across data centers later in this article.

This topology can be repeated across many data centers, if you configure SQL Server log shipping to one or more additional data centers.

Consult with your SAN vendor to determine whether you can use SAN replication or another supported mechanism to provide availability across data centers.

The following illustration shows primary and failover farms before failover.

**Primary and failover farms before failover**

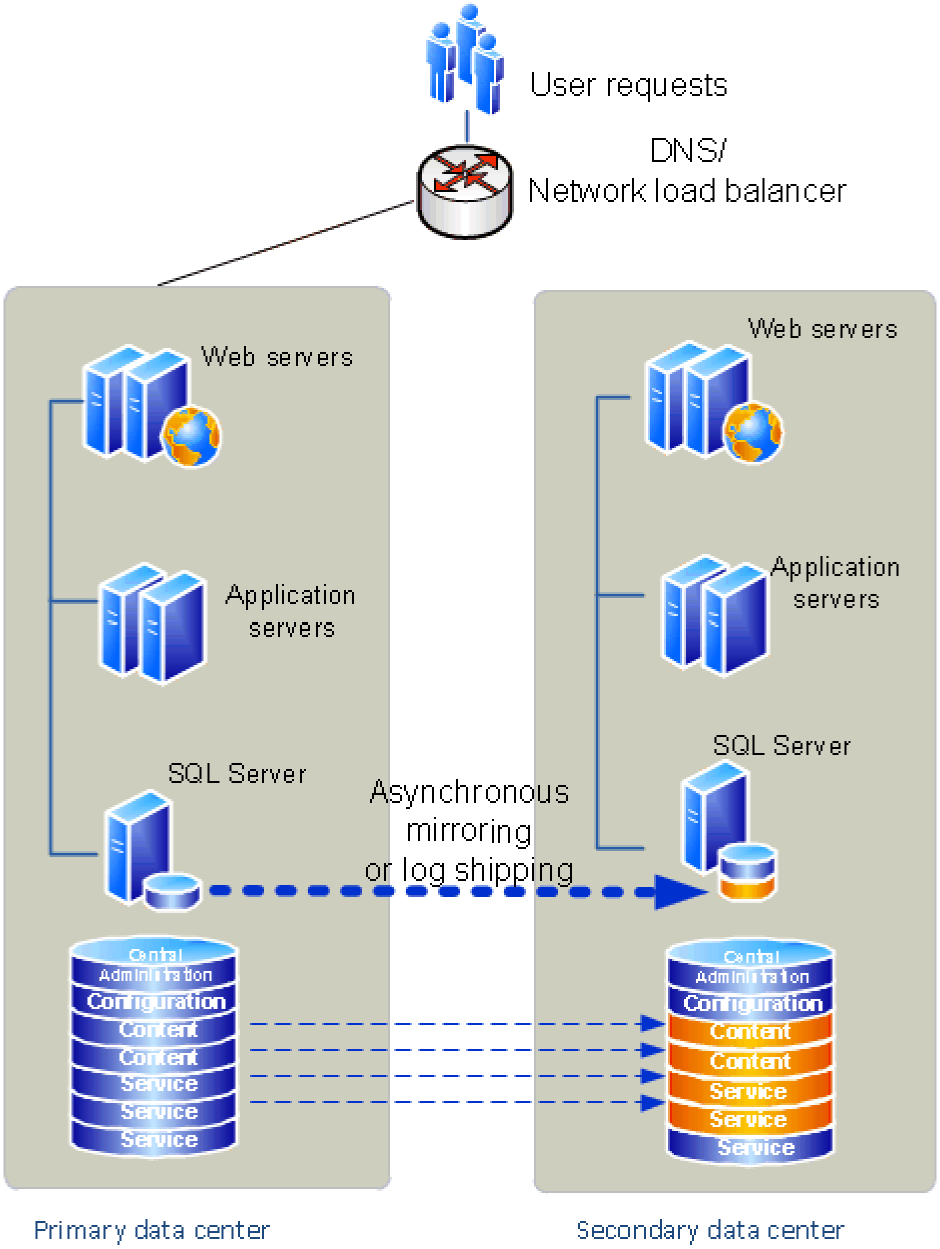

## Service application redundancy across data centers

To provide availability across data centers for service applications, we recommend that for the services that can be run cross-farm, you run a separate services farm that can be accessed from both the primary and the secondary data centers.

For services that cannot be run cross-farm, and to provide availability for the services farm itself, the strategy for providing redundancy across data centers for a service application varies. The strategy employed depends on whether:

- There is business value in running the service application in the disaster recovery farm when it is not in use.
- The databases associated with the service application can be log-shipped or asynchronously mirrored.
- The service application can run against read-only databases.

The following sections describe the disaster recovery strategies that we recommend for each service application. The service applications are grouped by strategy.

### Databases that can be log-shipped or asynchronously mirrored

After a service application has been initially deployed on a secondary farm, the databases that support the following service applications can be asynchronously mirrored or log-shipped across farms:

- **Application Registry service application**

  Databases: Application Registry service
- **Business Data Connectivity service application**

  Databases: Business Data Connectivity
- **Managed Metadata service application**

  Databases: Managed Metadata service

  **Note:**
  If tagging is in use, to successfully use the Managed Metadata service application in the disaster recovery farm, you must also log-ship or mirror the Tagging database for the User Profile service application.
- **PerformancePoint Services**

  Databases: PerformancePoint Service application
- **Project Server service application**

  Databases: Draft, Published, Archive, Reporting

  Project Server 2010 requires synchronization between its databases. Project Server can be replicated between farms by using an asynchronous replication mechanism (asynchronous database mirroring, log shipping, or asynchronous SAN replication), but, for recovery, you must ensure that the Project database logs are synchronized as you restore.

**Note:**
Although we recommend that you log-ship or mirror the Project Server databases to the disaster recovery farm, the Project Server service application cannot run against read-only databases. Therefore, we recommend that you do not run the Project Server service application on the disaster recovery farm until after failover. To successfully synchronize the Project Server databases on the disaster recovery farm, you must configure either time stamps or log marking for the databases.

- **Secure Store service application**

  Databases: Secure Store

- **Usage and Health Data Collection service application**

  Databases: Logging

  **Note:**
  It is possible to log-ship or mirror the Logging database. However, we recommend that you do not run the Usage and Health Data Collection service on the disaster recovery farm, and that you do not mirror nor log-ship the Logging database.

- **User Profile service application**

  Databases: Profile, Synchronization, Social Tagging

  The User Profile service Social Tagging database can be log-shipped. The Profile and Synchronization databases cannot be log-shipped.

  To provide redundancy for the User Profile service application, you must first deploy the service application in both the primary and secondary data centers.

  For the Social Tagging database, set up log-shipping.

  To set up the Profile and Synchronization databases, we recommend that you recover a backup of the databases to the secondary data center and attach them to the User Profile service application in that data center.

  To keep the profiles synchronized, you must run the User Profile Replication Engine that is included in the SharePoint Administration Toolkit after profile data has been updated on the primary farm. For more information, see User Profile Replication Engine overview (SharePoint Server 2010).

- **Web Analytics service application**

  Databases: Staging, Reporting

  **Note:**
  We recommend that you log-ship or mirror the Web Analytics Staging and Reporting databases. However, we recommend that you not run the Web Analytics service application on the disaster recovery farm until after failover.

## Service applications and databases that cannot be log-shipped or asynchronously mirrored

The following service applications must be deployed on both the primary and failover farms, and cannot be log-shipped or asynchronously mirrored. For most of these service applications, we recommend that you deploy them and then verify that the failover farm has the same configuration settings as the primary farm. If configuration changes that affect the service are made on the primary farm, you must update the failover farm.

- **Microsoft SharePoint Foundation Subscription Settings service application**

  Database: Subscription

  **Note:**
  Log-shipping the Subscription Settings database is not supported.

- **Access Services**

  Databases: None

- **Excel Services**

  Databases: None

- **Search**

  Databases: Crawl, Property, Search Administration

  Search requires complete synchronization between its databases and index. Because of this requirement, search cannot be replicated between farms by using an asynchronous replication mechanism (asynchronous database mirroring, log shipping, or asynchronous SAN replication).

  To provide up-to-date search on a failover farm, you must run search on the secondary farm.

  **Important:**
  The Search service application on the failover farm must be set to actively crawl the secondary farm. On failover, you must configure the Web application association to use the failover Search service application.

- **State service**

  Databases: State

  **Note:**
  Log-shipping the State database is not supported.

- **Visio Services**

  Databases: None

- **Word Automation Services**

  Databases: Word Automation Services

  Log-shipping the Word Automation Services database is not supported.

# System requirements for disaster recovery

In an ideal scenario, the failover components and systems match the primary components and systems in all ways: platform, hardware, and number of servers. At a minimum, the failover environment must be able to handle the traffic that you expect during a failover. Keep in mind that only a subset of users may be served by the failover site. The systems must match in at least the following:

- Operating system version and all updates
- SQL Server versions and all updates
- SharePoint 2010 Products versions and all updates

Although this article primarily discusses the availability of SharePoint 2010 Products, the system uptime will also be affected by the other components in the system. In particular, make sure that you do the following:

- Ensure that infrastructure dependencies such as power, cooling, network, directory, and SMTP are fully redundant.
- Choose a switching mechanism, whether DNS or hardware load balancing, that meets your needs.

# Global deployment of multiple farms (SharePoint Server 2010)

This section contains resources to help you design architectures that span geographic locations.

In this section:

- Global solutions for SharePoint 2010 Products (model)
- Client solutions for WAN environments (SharePoint Server 2010)

# Global solutions for SharePoint 2010 Products (model)

This model illustrates supported architectures for deploying Microsoft SharePoint 2010 Products geographically.

Global Solutions for SharePoint 2010 Products

Visio (http://go.microsoft.com/fwlink/?LinkId=206424)

PDF (http://go.microsoft.com/fwlink/?LinkId=206429)

XPS (http://go.microsoft.com/fwlink/?LinkId=206432)

# Client solutions for WAN environments (SharePoint Server 2010)

Many organizations include users who are distributed across WAN links and users who are not always connected to the network. Microsoft SharePoint 2010 Products accommodate various network scenarios and work environments. This article describes the client solutions that can be used across slow network connections or when users are offline.

In this article:

- Mobile views
- <token xmlns="http://ddue.schemas.microsoft.com/authoring/2003/5">OfficeWebAccess_2nd_CurrentVer</token>
- <token xmlns="http://ddue.schemas.microsoft.com/authoring/2003/5">Office_2nd_CurrentVer</token> Document Cache and the MS-FSSHTTP protocol
- <token xmlns="http://ddue.schemas.microsoft.com/authoring/2003/5">Outlook_2nd_CurrentVer</token>
- <token xmlns="http://ddue.schemas.microsoft.com/authoring/2003/5">Groove_2nd_NoVer</token>
- SharePoint Workspace Mobile for Windows Phone 7
- <token xmlns="http://ddue.schemas.microsoft.com/authoring/2003/5">Groove_2nd_NoVer</token> with <token xmlns="http://ddue.schemas.microsoft.com/authoring/2003/5">Groove_Server_short_current_NoVer</token>

The following tables compare the solutions by indicating the circumstances in which a solution works best and the scope of access a solution provides for SharePoint sites. The rest of this article describes the solutions in more detail.

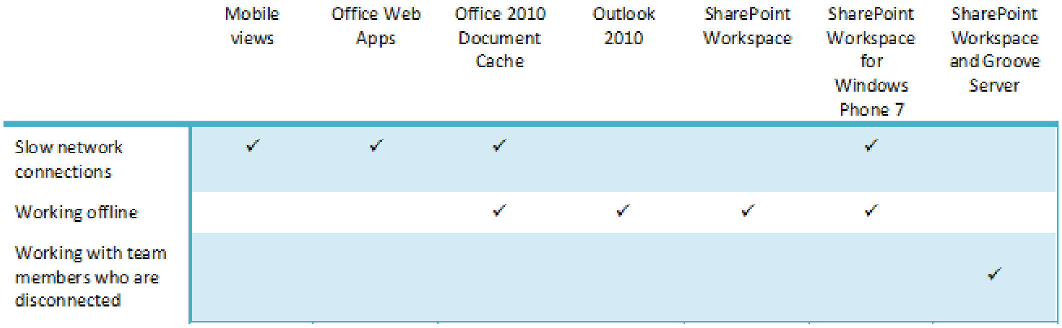

| | Mobile views | Office Web Apps | Office 2010 Document Cache | Outlook 2010 | SharePoint Workspace | SharePoint Workspace for Windows Phone 7 | SharePoint Workspace and Groove Server |
|---|---|---|---|---|---|---|---|
| Slow network connections | ✓ | ✓ | ✓ | | | ✓ | |
| Working offline | | | ✓ | ✓ | ✓ | ✓ | |
| Working with team members who are disconnected | | | | | | | ✓ |

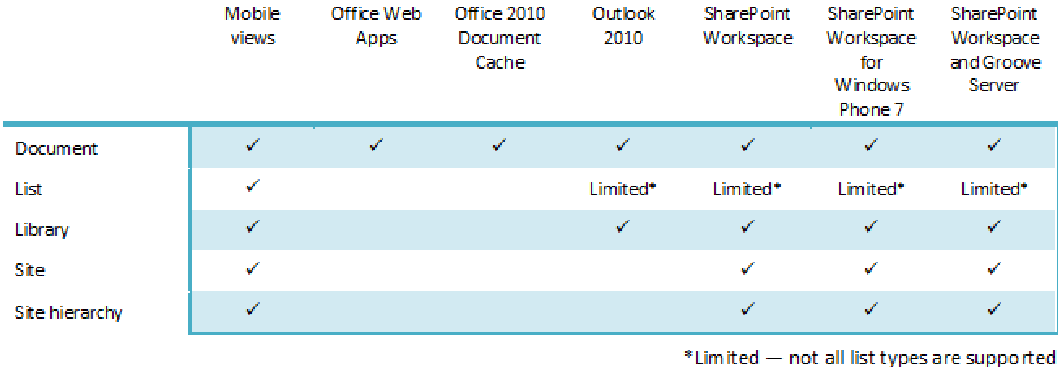

| | Mobile views | Office Web Apps | Office 2010 Document Cache | Outlook 2010 | SharePoint Workspace | SharePoint Workspace for Windows Phone 7 | SharePoint Workspace and Groove Server |
|---|---|---|---|---|---|---|---|
| Document | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| List | ✓ | | | Limited* | Limited* | Limited* | Limited* |
| Library | ✓ | | | ✓ | ✓ | ✓ | ✓ |
| Site | ✓ | | | | ✓ | ✓ | ✓ |
| Site hierarchy | ✓ | | | | ✓ | ✓ | ✓ |

*Limited — not all list types are supported

# Mobile views

Mobile views are not only for mobile devices. In low bandwidth or high-latency environments — such as latencies of 300 milliseconds or more — mobile views can provide acceptable performance when they are used to navigate a site hierarchy, complete simple forms, and view textual data.

To display a mobile view in a non-mobile browser, append **?mobile=1** to the URL of any SharePoint site.

By default, mobile views are enabled for most lists or libraries that are created by using default templates. By default, mobile views are not enabled for custom lists, custom libraries, or for lists or libraries that were created in previous versions of the product that have been upgraded to Microsoft SharePoint Foundation 2010. Mobile views are not available for Datasheet and Gantt view types.

To configure mobile views for lists and libraries, see Configure mobile views (SharePoint Server 2010).

# Office Web Apps

Microsoft Office Web Apps are online companions to Microsoft Word, Microsoft Excel, Microsoft PowerPoint, and Microsoft OneNote, which enable people to access files and do light editing in a browser without downloading and uploading files over low-bandwidth or high-latency connections.

When Office Web Apps are enabled for a site collection, you can choose to view and edit files in a browser, as shown in the following illustration.

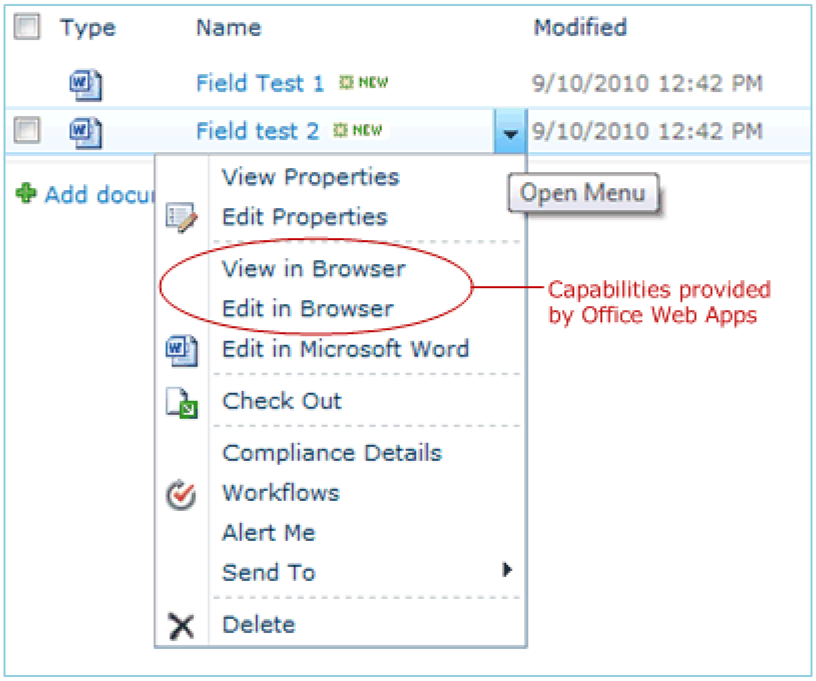

| Type | Name | Modified |
|---|---|---|
|  | Field Test 1 NEW | 9/10/2010 12:42 PM |
|  | Field test 2 NEW | 9/10/2010 12:42 PM |

Capabilities provided by Office Web Apps

In most cases, opening files in a browser results in a faster time-to-first-page than opening files in one of the Microsoft Office 2010 client applications.

Additionally, users in your organization can use browser-enabled cell phones and mobile devices to read Word, Excel, and PowerPoint documents that are stored on a SharePoint Server computer if views and content that are enabled for mobile access are published outside a firewall. The following devices provide mobile support for Office Web Apps:

- Windows Phone 7
- Windows Mobile
- BlackBerry
- iPhone, iPod Touch
- Nokia S60
- Feature telephones from Japan, including NTT DOCOMO, SoftBank, and au by KDDI

Microsoft Silverlight can improve the Office Web Apps user experience. Silverlight is a free plug-in that can provide richer Web experiences for many browsers. The Silverlight plug-in is not required to be installed on the client browser to use Office Web Apps. However, installing the Silverlight plug-in on the browser can provide the following benefits:

- When users use Word Web App on browsers that have the Silverlight plug-in installed, they can experience faster page loading, improved text fidelity at full zoom, Microsoft ClearType tuner settings support, and improved accuracy in location of search string instances when they use the Find on this Page command.
- When users use PowerPoint Web App on browsers that have the Silverlight plug-in installed, they can experience faster page loading, animations will appear smoother, and presentation slides will scale with the browser window size.

There are no additional benefits in Excel Web App and OneNote Web App when Silverlight is installed on the client browser.

For more information about Silverlight, see Microsoft Silverlight (http://go.microsoft.com/fwlink/?LinkId=206068).

For more information about Office Web Apps, see Microsoft Web Apps Deployment (http://go.microsoft.com/fwlink/?LinkId=206018).

# Office 2010 Document Cache and the MS-FSSHTTP protocol

Office 2010 combined with SharePoint 2010 Products improves the experience of users who are using and managing files over slow network connections. Files that are open from SharePoint 2010 sites are downloaded by using asynchronous file transfer and cached locally. As a result, files open more quickly and users can start using the file before the download is complete. Additionally, only changes to files — not whole files — are transmitted between SharePoint 2010 Products and client computers. These capabilities are made possible by the new File Synchronization via SOAP over HTTP (MS-FSSHTTP) protocol.

The Document Cache settings and features can be accessed and managed through the Upload Center, a new feature of Office 2010 that is automatically installed. Users can access the Upload Center by either clicking the Upload Center icon in the notification area or by opening it from the Start menu.

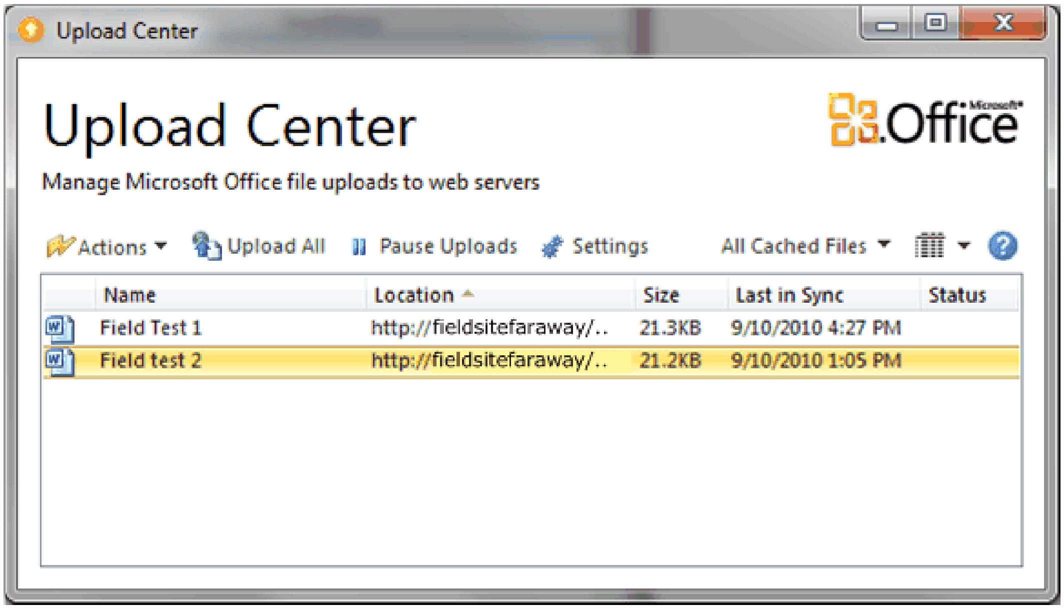

Upload Center

Manage Microsoft Office file uploads to web servers

| Name | Location | Size | Last in Sync | Status |
|---|---|---|---|---|
| Field Test 1 | http://fieldsitefaraway/.. | 21.3KB | 9/10/2010 4:27 PM | |
| Field test 2 | http://fieldsitefaraway/.. | 21.2KB | 9/10/2010 1:05 PM | |

The Upload Center lists all files that have been cached. Cached versions of files can be taken offline and managed through the Upload Center. Users can monitor the status of files that are in the process of being uploaded. As shown in the following illustration, users can also manage cache settings to determine how long cached files are retained and to delete all cached files, if necessary.

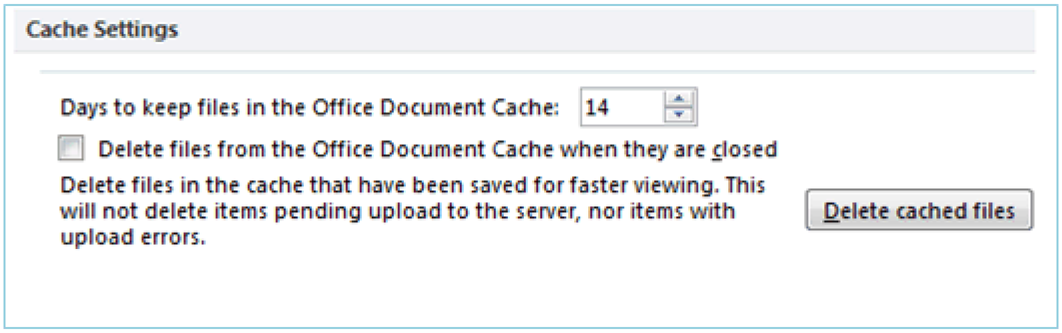

Cache Settings

Days to keep files in the Office Document Cache: 14

□ Delete files from the Office Document Cache when they are closed

Delete files in the cache that have been saved for faster viewing. This will not delete items pending upload to the server, nor items with upload errors.

Users do not have to access and manage the Upload Center to take advantage of the caching abilities of Office 2010. The functionality of Office 2010 takes place in the background, even without user intervention.

**Note:**

When users are working with earlier versions of Office, such as Office 2007, BranchCache can be used to reduce bandwidth utilization and download times for frequently accessed content. BranchCache, when it is enabled on the server and client computers, greatly reduces the time to open files that have been cached. However, BranchCache does not reduce the amount of network bandwidth that is used when a file is saved back to a SharePoint site. BranchCache cannot be used in combination with Document Cache and the MS-FSSHTTP protocol. If both BranchCache and Document Cache are implemented, communication takes place by using the MS-FSSHTTP protocol. The use of SharePoint Server 2010 in combination with Office 2010 results in significantly better WAN performance overall.

For more information about how to use BrancheCache, see BranchCache in Windows 7 and Windows Server 2008 R2 Overview (http://go.microsoft.com/fwlink/?LinkId=206069).

For more information about the Upload Center and Office Document Cache settings, see the following articles:

- Microsoft Office 2010 Upload Center (http://go.microsoft.com/fwlink/?LinkId=206070)
- Office Document Cache settings (http://go.microsoft.com/fwlink/?LinkId=206071)

# Outlook 2010

Users can synchronize a SharePoint library, contact list, task list, project task list, and a certain type of SharePoint external list with Outlook 2010. Because many users of SharePoint 2010 Products also use Outlook 2010 to collaborate and coordinate activities and projects, the ability to synchronize these libraries and lists can help users become more efficient, especially when they are working offline or when access to SharePoint sites is not convenient.

Just as in other Office 2010 applications, SharePoint files that are synchronized with Outlook 2010 are downloaded, uploaded, and cached by using the MS-FSSHTTP protocol. The result is better performance over slow network connections than the performance in previous versions.

To synchronize a SharePoint library with Outlook 2010, on the **Library** tab, in the **Connect & Export** group, click **Connect to Outlook**, as shown in the following illustration.

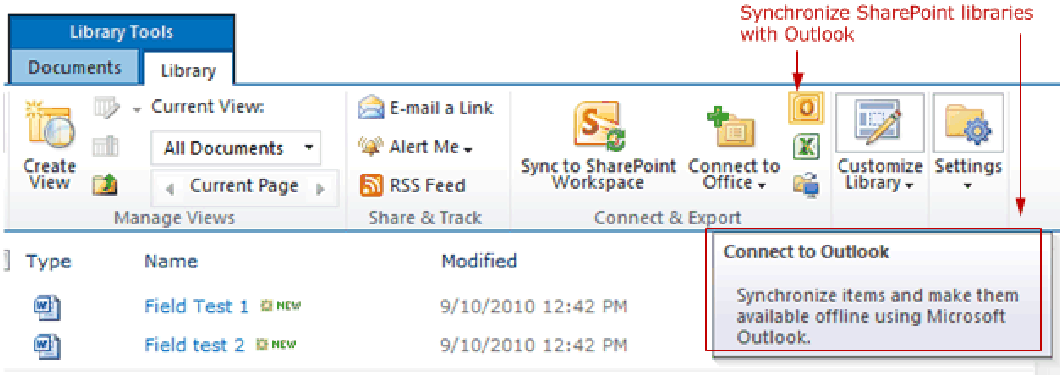
Synchronize SharePoint libraries with Outlook

Connect to Outlook

Synchronize items and make them available offline using Microsoft Outlook.

| Type | Name | Modified |
|---|---|---|
| | Field Test 1 NEW | 9/10/2010 12:42 PM |
| | Field test 2 NEW | 9/10/2010 12:42 PM |

When users are working offline, changes to files are saved locally. When they are online again, users can select to either upload changes to the SharePoint library when the client application first synchronizes or to keep the changed files locally.

Consider the following when planning to use Outlook 2010 in a WAN environment:

- Outlook 2010 does not provide peer-to-peer synchronization of content in SharePoint sites. Outlook 2010 provides a personal offline experience for individual team members.
- By default, synchronized files are not checked out automatically. A user cannot check out a file within Outlook 2010. The recommended practice is to check out files within SharePoint 2010 Products before synchronizing and editing the content in Outlook 2010.

For more information, see Synchronize SharePoint 2010 content with Outlook 2010 (http://go.microsoft.com/fwlink/?LinkId=206097).

# SharePoint Workspace

Microsoft SharePoint Workspace 2010, formerly Microsoft Office Groove, is a client application for Microsoft SharePoint Server 2010 and Microsoft SharePoint Foundation 2010 that supports online and offline collaboration.

SharePoint Workspace 2010 builds on the previous version and adds powerful new tools that let users access and share content that is stored on SharePoint sites, even when users are not connected to corporate networks. SharePoint Workspace 2010 is included with Microsoft Office Professional Plus 2010 and provides the most robust offline experience for SharePoint sites, allowing users to take entire SharePoint sites offline.

The following illustration demonstrates how to synchronize a site with SharePoint Workspace 2010.

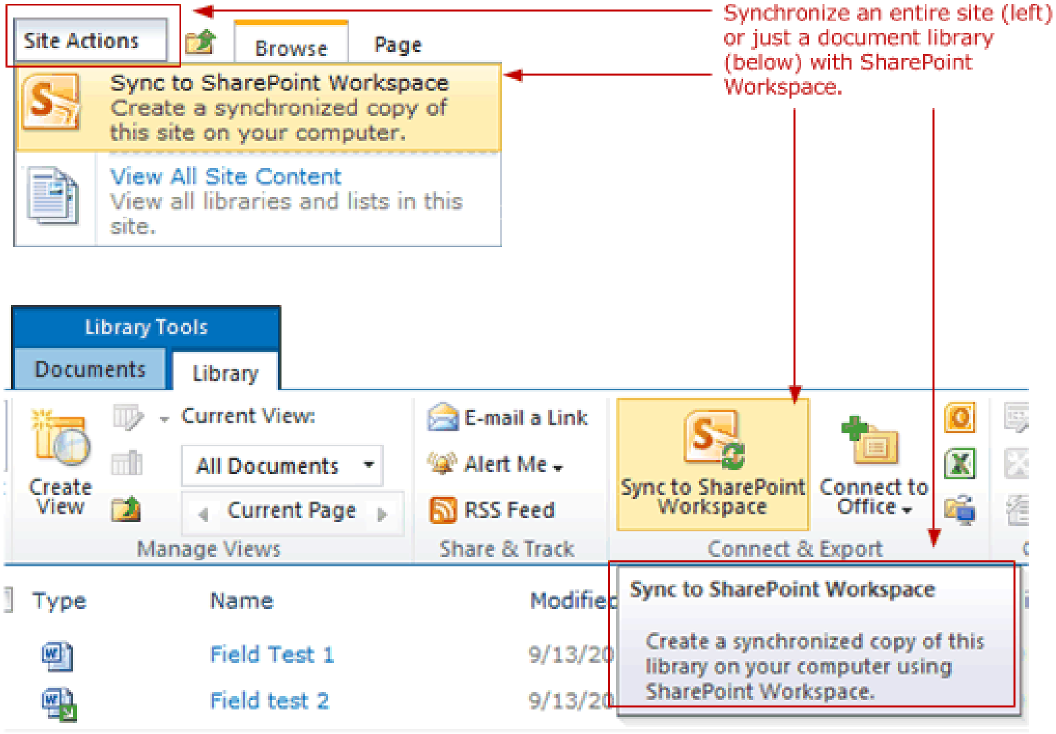

Synchronize an entire site (left) or just a document library (below) with SharePoint Workspace.

Sync to SharePoint Workspace
Create a synchronized copy of this site on your computer.

View All Site Content
View all libraries and lists in this site.

Sync to SharePoint Workspace

Create a synchronized copy of this library on your computer using SharePoint Workspace.

| Type | Name | Modified |
|---|---|---|
| | Field Test 1 | 9/13/20 |
| | Field test 2 | 9/13/20 |

For more information, see What's new in SharePoint Workspace 2010 (http://go.microsoft.com/fwlink/?LinkId=206098).

# SharePoint Workspace Mobile for Windows Phone 7

Windows Phone 7 includes Microsoft Office Mobile, which lets users work with files from their telephones. Microsoft SharePoint Workspace Mobile is part of Office Mobile and already on the telephone in the Office Hub, as shown in the following illustration.

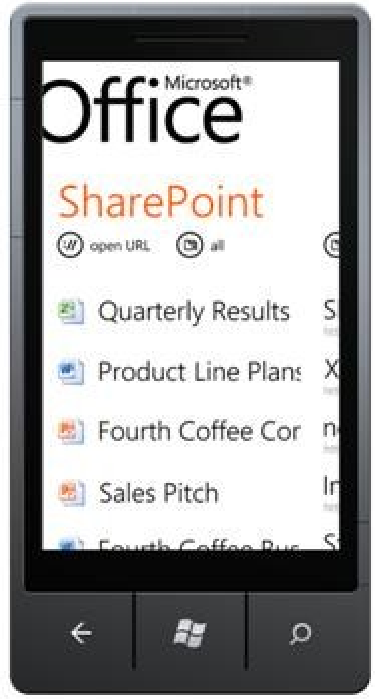

Mobile users can use Windows Phone 7 and Microsoft SharePoint Workspace Mobile to do the following:

- View content hosted on a SharePoint 2010 site.
- Open and edit Word, Excel, PowerPoint, and OneNote files that are hosted on a SharePoint 2010 site.
- Browse SharePoint 2010 sites, lists, and document libraries.
- Increase the security of remote access to corporate resources through Microsoft Forefront Unified Access Gateway (UAG), if their company uses it.

Users can also use Microsoft SharePoint Workspace Mobile to take SharePoint 2010 files offline on the telephone. Users can open and edit the files, and then save them back to the SharePoint site when the users are back online.

For more information, see Office Mobile 2010 for Windows Phone 7 (http://go.microsoft.com/fwlink/?LinkId=206099).

# SharePoint Workspace with Groove Server

SharePoint Workspace 2010 can also be used together with Microsoft Groove Server to provide peer collaboration that does not require all team members to be connected to SharePoint 2010 Products. Collaboration can be extended outside a private network to trusted partners and field sites. With this configuration, at least one team member must have access to the SharePoint sites. Other team members must have access to either the computer of another team member who has access to the SharePoint sites, or access directly to the SharePoint sites. The following illustration provides an example of this architecture.

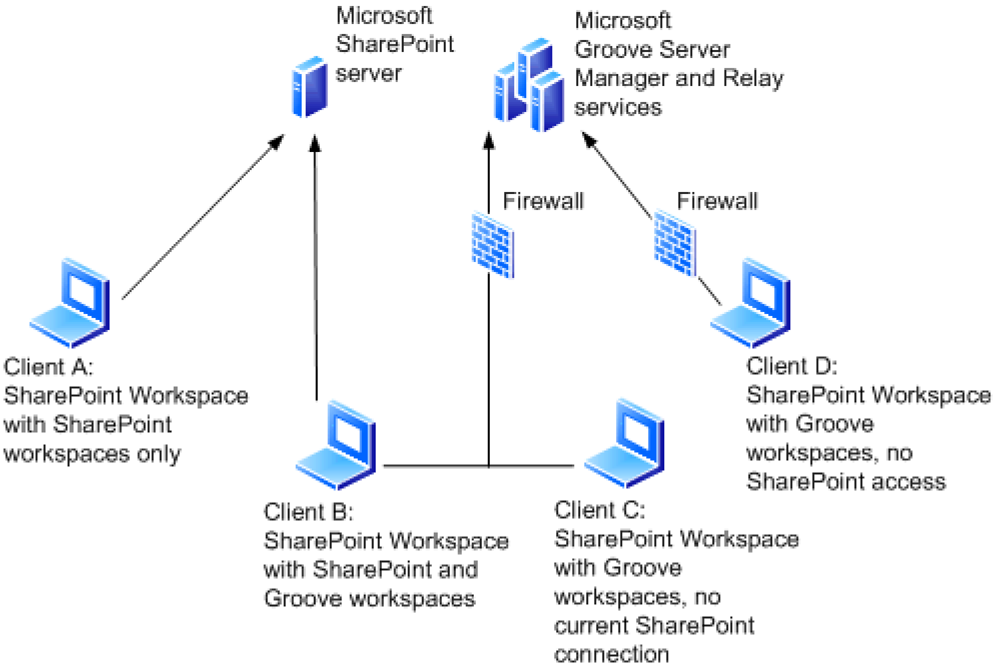
Microsoft SharePoint server
Microsoft Groove Server Manager and Relay services
Firewall
Firewall
Client A: SharePoint Workspace with SharePoint workspaces only
Client B: SharePoint Workspace with SharePoint and Groove workspaces
Client C: SharePoint Workspace with Groove workspaces, no current SharePoint connection
Client D: SharePoint Workspace with Groove workspaces, no SharePoint access

For more information about SharePoint Workspace and Groove Server, see Plan for SharePoint Workspace 2010.

# Planning worksheets for SharePoint Server 2010

In this article:

- Planning worksheets by task
- Planning worksheets by title

This article provides links to worksheets that you can use to record information that you gather and decisions that you make as you plan your deployment of Microsoft SharePoint Server 2010. Use these worksheets in conjunction with — not as a substitute for — Planning and architecture for SharePoint Server 2010.

## Planning worksheets by task

| For this task | Use this worksheet | To do this |
|---|---|---|
| Plan sites and site collections (SharePoint Server 2010) | Site planning data worksheet (http://go.microsoft.com/fwlink/?LinkID=167837) | Plan top level site collections and sites, and record decisions about site themes and navigation. |
| Plan site navigation (SharePoint Server 2010) | Site planning data worksheet (http://go.microsoft.com/fwlink/?LinkID=167837) | Plan top level site collections and sites, and record decisions about site themes and navigation. |
| Plan for using themes (SharePoint Server 2010) | Site planning data worksheet (http://go.microsoft.com/fwlink/?LinkID=167837) | Plan top level site collections and sites, and record decisions about site themes and navigation. |
| Plan incoming e-mail (SharePoint Server 2010) | Plan incoming e-mail worksheet (http://go.microsoft.com/fwlink/?LinkId=200542) | Plan incoming e-mail in order to enable SharePoint |

| For this task | Use this worksheet | To do this |
|---|---|---|
| | | sites to receive and store e-mail messages and attachments in lists and libraries. |
| Plan content deployment (SharePoint Server 2010) | Content deployment data worksheet (http://go.microsoft.com/fwlink/?LinkID=167835) | Plan the export and import servers in the farms in your content deployment topology, and to plan the content deployment paths and jobs. |
| Plan managed metadata (SharePoint Server 2010) | Term sets planning worksheet(http://go.microsoft.com/fwlink/?LinkId=163486) | Determine basic taxonomy, including term, usage, owner, and group. |
| Plan managed metadata (SharePoint Server 2010) | Detailed term set planning worksheet(http://go.microsoft.com/fwlink/?LinkId=163487) | Determine taxonomy including detailed identifying characteristics such as measurements. |
| Plan managed metadata (SharePoint Server 2010) | Managed metadata services planning worksheet(http://go.microsoft.com/fwlink/?LinkId=164578) | Plan to share metadata information using managed metadata services and connections. |
| Document management planning (SharePoint Server 2010) | Document management participants worksheet (http://go.microsoft.com/fwlink/?LinkID=165871) | Identify document management planning stakeholders and record document management practices. |
| Document management planning | Analyze document usage worksheet (http://go.microsoft.com/fwlink/?LinkID=165873) | Record information gathered when analyzing document |

| For this task | Use this worksheet | To do this |
|---|---|---|
| (SharePoint Server 2010) | | usage. |
| Document management planning (SharePoint Server 2010) | Policy worksheet (http://go.microsoft.com/fwlink/?LinkID=165883) | Plan information management policies for content types. |
| Records management planning (SharePoint Server 2010) | In-place records planning worksheet(http://go.microsoft.com/fwlink/?LinkId=185011) | Identify record types and content types to be stored in normal document libraries. |
| Plan for backup and recovery (SharePoint Server 2010) | Backup and recovery planning workbook (http://go.microsoft.com/fwlink/?LinkID=184385) | Help you plan strategies for backup and recovery for SharePoint Server 2010 environment. |
| Document management planning (SharePoint Server 2010) | Document management planning (SharePoint Server 2010) | Plan a content type. |
| Plan and prepare for upgrade (SharePoint Server 2010) | Upgrade worksheet (http://go.microsoft.com/fwlink/?LinkId=179928) | Record information about your environment while you prepare for upgrade. |
| Metadata-based routing and storage planning (SharePoint Server 2010) | Content Organizer settings worksheet (http://go.microsoft.com/fwlink/?LinkId=189018&clcid=0x409) | Determine and record how the content organizer settings in your site can be an effective part of your metadata-based content routing and storage solution. |

| For this task | Use this worksheet | To do this |
|---|---|---|
| Metadata-based routing and storage planning (SharePoint Server 2010) | Content Organizer rule worksheet (http://go.microsoft.com/fwlink/?LinkId=189019&clcid=0x409) | Plan rules that will be an effective part of your metadata-based routing and storage solution. |

## Planning worksheets by title

| Use this worksheet | For this task | To do this |
|---|---|---|
| Analyze document usage worksheet (http://go.microsoft.com/fwlink/?LinkID=165873) | Document management planning (SharePoint Server 2010) | Record information gathered when analyzing document usage. |
| Backup and recovery planning workbook (http://go.microsoft.com/fwlink/?LinkID=184385) | Plan for backup and recovery (SharePoint Server 2010) | Help you plan strategies for backup and recovery for SharePoint Server 2010 environment. |
| Content deployment data worksheet (http://go.microsoft.com/fwlink/?LinkID=167835) | Plan content deployment (SharePoint Server 2010) | Plan the export and import servers in the farms in your content deployment topology, and to plan the content deployment paths and jobs. |
| Content Organizer rules worksheet (http://go.microsoft.com/fwlink/?LinkId=189019&clcid=0x409) | Metadata-based | Plan rules that will be an |

| Use this worksheet | For this task | To do this |
|---|---|---|
|  | routing and storage planning (SharePoint Server 2010) | effective part of your metadata-based routing and storage solution. |
| Content Organizer settings worksheet (http://go.microsoft.com/fwlink/?LinkID=167835) | Metadata-based routing and storage planning (SharePoint Server 2010) | Determine and record how the content organizer settings in your site can be an effective part of your metadata-based content routing and storage solution. |
| Content type worksheet (http://go.microsoft.com/fwlink/?LinkID=165878) | Document management planning (SharePoint Server 2010) | Plan a content type. |
| Detailed term set planning worksheet(http://go.microsoft.com/fwlink/?LinkId=163487&clcid=0x409 ) | Plan managed metadata (SharePoint Server 2010) | Determine taxonomy including detailed identifying characteristics such as measurements . |
| Document libraries worksheet (http://go.microsoft.com/fwlink/?LinkID=165874) | Document management planning (SharePoint Server 2010) | Plan libraries based on sites and on document types. |
| Document management participants worksheet (http://go.microsoft.com/fwlink/?LinkID=165871) | Document managemen | Identify document |

| Use this worksheet | For this task | To do this |
|---|---|---|
| | t planning (SharePoint Server 2010) | management planning stakeholders and record document management practices. |
| In-place records planning worksheet(http://go.microsoft.com/fwlink/?LinkId=185011&clcid=0x409) | Records management planning (SharePoint Server 2010) | Identify record types and content types to be stored in normal document libraries. |
| Managed metadata services planning worksheet(http://go.microsoft.com/fwlink/?LinkId=164578) | Plan managed metadata (SharePoint Server 2010) | Plan to share metadata information using managed metadata services and connections. |
| Plan incoming e-mail worksheet (http://go.microsoft.com/fwlink/?LinkId=200542) | Plan incoming e-mail (SharePoint Server 2010) | Plan incoming e-mail in order to enable SharePoint sites to receive and store e-mail messages and attachments in lists and libraries. |
| Policy worksheet (http://go.microsoft.com/fwlink/?LinkID=165883) | Document management planning (SharePoint Server 2010) | Plan information management policies for content types. |
| Site planning data worksheet | Plan sites | Plan top level |

| Use this worksheet | For this task | To do this |
|---|---|---|
| (http://go.microsoft.com/fwlink/?LinkID=167837) | and site collections (SharePoint Server 2010)<br>Plan site navigation (SharePoint Server 2010)<br>Plan for using themes (SharePoint Server 2010) | site collections and sites, and record decisions about site themes and navigation. |
| Term sets planning worksheet(http://go.microsoft.com/fwlink/?LinkId=163486) | Plan managed metadata (SharePoint Server 2010) | Determine basic taxonomy, including term, usage, owner, and group. |
| Upgrade worksheet (http://go.microsoft.com/fwlink/?LinkId=179928) | Plan and prepare for upgrade (SharePoint Server 2010) | Record information about your environment while you prepare for upgrade. |