[MS-WSSO]:

Windows SharePoint Services Overview

Intellectual Property Rights Notice for Open Specifications Documentation

- **Technical Documentation.** Microsoft publishes Open Specifications documentation ("this documentation") for protocols, file formats, data portability, computer languages, and standards support. Additionally, overview documents cover inter-protocol relationships and interactions.
- **Copyrights**. This documentation is covered by Microsoft copyrights. Regardless of any other terms that are contained in the terms of use for the Microsoft website that hosts this documentation, you can make copies of it in order to develop implementations of the technologies that are described in this documentation and can distribute portions of it in your implementations that use these technologies or in your documentation as necessary to properly document the implementation. You can also distribute in your implementation, with or without modification, any schemas, IDLs, or code samples that are included in the documentation. This permission also applies to any documents that are referenced in the Open Specifications documentation.
- **No Trade Secrets**. Microsoft does not claim any trade secret rights in this documentation.
- **Patents**. Microsoft has patents that might cover your implementations of the technologies described in the Open Specifications documentation. Neither this notice nor Microsoft's delivery of this documentation grants any licenses under those patents or any other Microsoft patents. However, a given Open Specifications document might be covered by the Microsoft Open Specifications Promise or the Microsoft Community Promise. If you would prefer a written license, or if the technologies described in this documentation are not covered by the Open Specifications Promise or Community Promise, as applicable, patent licenses are available by contacting iplg@microsoft.com.
- **Trademarks**. The names of companies and products contained in this documentation might be covered by trademarks or similar intellectual property rights. This notice does not grant any licenses under those rights. For a list of Microsoft trademarks, visit www.microsoft.com/trademarks.
- **Fictitious Names**. The example companies, organizations, products, domain names, email addresses, logos, people, places, and events that are depicted in this documentation are fictitious. No association with any real company, organization, product, domain name, email address, logo, person, place, or event is intended or should be inferred.

Reservation of Rights. All other rights are reserved, and this notice does not grant any rights other than as specifically described above, whether by implication, estoppel, or otherwise.

Tools. The Open Specifications documentation does not require the use of Microsoft programming tools or programming environments in order for you to develop an implementation. If you have access to Microsoft programming tools and environments, you are free to take advantage of them. Certain Open Specifications documents are intended for use in conjunction with publicly available standards specifications and network programming art and, as such, assume that the reader either is familiar with the aforementioned material or has immediate access to it.

Revision Summary

Date	Revision History	Revision Class	Comments
3/14/2008	0.1	Major	Initial Availability.
6/20/2008	0.1.1	Editorial	Revised and edited the technical content.
7/25/2008	0.1.2	Editorial	Revised and edited the technical content.
8/29/2008	1.0	Major	Updated and revised the technical content.
10/24/2008	1.0.1	Editorial	Revised and edited the technical content.
12/5/2008	1.0.2	Editorial	Initial availability
1/16/2009	1.0.3	Editorial	Revised and edited the technical content.
2/27/2009	1.0.4	Editorial	Revised and edited the technical content.
4/10/2009	2.0	Major	Updated and revised the technical content.
5/22/2009	2.0.1	Editorial	Revised and edited the technical content.
7/2/2009	3.0	Major	Updated and revised the technical content.
8/14/2009	3.0.1	Editorial	Revised and edited the technical content.
9/25/2009	3.0.2	Editorial	Revised and edited the technical content.
11/6/2009	3.0.3	Editorial	Revised and edited the technical content.
12/18/2009	4.0	Major	Updated and revised the technical content.
1/29/2010	4.0.1	Editorial	Revised and edited the technical content.
3/12/2010	5.0	Major	Updated and revised the technical content.
4/23/2010	6.0	Major	Updated and revised the technical content.
6/4/2010	6.0.1	Editorial	Revised and edited the technical content.
7/16/2010	6.0.1	None	No changes to the meaning, language, or formatting of the technical content.
8/27/2010	6.0.1	None	No changes to the meaning, language, or formatting of the technical content.
10/8/2010	7.0	Major	Significantly changed the technical content.
11/19/2010	7.1	Minor	Clarified the meaning of the technical content.
1/7/2011	7.1	None	No changes to the meaning, language, or formatting of the technical content.
2/11/2011	7.1	None	No changes to the meaning, language, or formatting of the technical content.
3/25/2011	7.1	None	No changes to the meaning, language, or formatting of the technical content.
5/6/2011	7.1	None	No changes to the meaning, language, or formatting of the technical content.

Date	Revision History	Revision Class	Comments
6/17/2011	7.2	Minor	Clarified the meaning of the technical content.
9/23/2011	7.3	Minor	Clarified the meaning of the technical content.
12/16/2011	7.4	Minor	Clarified the meaning of the technical content.
3/30/2012	7.4	None	No changes to the meaning, language, or formatting of the technical content.
7/12/2012	7.5	Minor	Clarified the meaning of the technical content.
10/25/2012	7.5	None	No changes to the meaning, language, or formatting of the technical content.
1/31/2013	7.5	None	No changes to the meaning, language, or formatting of the technical content.
8/8/2013	7.5	None	No changes to the meaning, language, or formatting of the technical content.
11/14/2013	7.5	None	No changes to the meaning, language, or formatting of the technical content.
4/30/2014	7.5	None	No changes to the meaning, language, or formatting of the technical content.
7/31/2014	7.5	None	No changes to the meaning, language, or formatting of the technical content.
10/30/2014	7.5	None	No changes to the meaning, language, or formatting of the technical content.
6/23/2016	7.5	None	No changes to the meaning, language, or formatting of the technical content.
9/14/2016	7.5	None	No changes to the meaning, language, or formatting of the technical content.

Table of Contents

1 Introduction

This document provides an informative overview of the back-end protocols implemented by Windows SharePoint Services File, Print, and User/Group administration capabilities. Windows SharePoint Services is a web-based technology that provides:

- A ready-to-use, team-oriented **website** for collaboration.
- A development platform for building web-based experiences that take advantage of the collaboration features of Windows SharePoint Services.
- A framework for deploying and managing the Windows SharePoint Services Team Site and applications built on the Windows SharePoint Services platform.

As part of the collaboration services, Windows SharePoint Services provides support for document collaboration, including the ability to store, update, and view documents. This capability is delivered through document libraries within Team Sites. Much of the Windows SharePoint Services infrastructure is designed to ensure that Windows SharePoint Services sites (2) provide these services in a highly scalable, manageable, and extensible way, as described in detail later in this document.

The purpose of this document is to provide an understanding of the concepts and architecture underlying the file management and security related features of Windows SharePoint Services. In order to deliver these file services capabilities, Windows SharePoint Services uses three major sets of protocols:

- File-oriented communication between the **end-user client** and the Windows SharePoint Services **front-end web server** using the **Web Distributed Authoring and Versioning Protocol (WebDAV)** as described in [RFC2518], [MS-WDV], and [MS-WDVSE]. The end-user client can also use the FrontPage Server Extensions Remote Protocol as described in [MS-FPSE]. The use of WebDAV is recommended over the FrontPage Server Extensions Remote Protocol.
- Web pages presented to the client using standard **Hypertext Transfer Protocol (HTTP)**.
- Communication between the front-end web server and the Windows SharePoint Services **back-end database server** using specific queries and stored procedures implemented using Tabular Data Stream (TDS), a protocol for SQL communication described in [MS-TDS]. Details of the File, Print, and User/Group administration communication between the front-end web server and back-end database server is described in [MS-WSSFOB] for Windows SharePoint Services 2.0, [MS-WSSFO] for Windows SharePoint Services 3.0, and [MS-WSSFO2] for Microsoft SharePoint Foundation 2010.

This document provides an overview for protocols for Windows SharePoint Services 2.0, Windows SharePoint Services 3.0, and SharePoint Foundation 2010. It generally refers to Windows SharePoint Services when the subject applies to all versions, and explicitly calls out the version when necessary for clarity.

Note This document will not be updated to reflect new releases. For more information, see the SharePoint Products and Technologies Protocols Overview [MS-SPO]. This document provides an informative overview of the front- and back-end protocols that are implemented by all SharePoint Products and Technologies for communicating with client and server applications. Unlike this overview, this document is not limited only to those protocols that provide file, print, and user/group administration capabilities. In addition, [MS-SPO] covers Windows SharePoint Services 3.0 and Microsoft Office SharePoint Server 2007 as well as SharePoint Foundation 2010 and Microsoft SharePoint Server 2010. Going forward, this document will also cover future versions of SharePoint Foundation 2010 and SharePoint Server 2010.

Note The **Transact-Structured Query Language (T-SQL)** based protocols change significantly in their function between Windows SharePoint Services 2.0, Windows SharePoint Services 3.0, and

SharePoint Foundation 2010. Separate versions of the specification documents for these protocols target each release.

- Windows SharePoint Services 2.0: [MS-WSSFOB] Windows SharePoint Services (Windows SharePoint Services): File Operations Database Communications Base Protocol Specification.
- Windows SharePoint Services 3.0: [MS-WSSFO] Windows SharePoint Services (Windows SharePoint Services): File Operations Database Communications Protocol Specification
- SharePoint Foundation 2010: [MS-WSSFO2] Windows SharePoint Services (Windows SharePoint Services): File Operations Database Communications Version 2 Protocol Specification

1.1 Glossary

This document uses the following terms:

access control entry (ACE): An entry in an **access control list (ACL)** that contains a set of user rights and a **security identifier (SID)** that identifies a principal for whom the rights are allowed, denied, or audited.

access control list (ACL): A list of **access control entries (ACEs)** that collectively describe the security rules for authorizing access to some resource; for example, an object or set of objects.

Active Directory: A general-purpose network **directory service**. **Active Directory** also refers to the Windows implementation of a **directory service**. **Active Directory** stores information about a variety of objects in the network. Importantly, user accounts, computer accounts, groups, and all related credential information used by the Windows implementation of **Kerberos** are stored in **Active Directory**. **Active Directory** is either deployed as Active Directory Domain Services (AD DS) or Active Directory Lightweight Directory Services (AD LDS). [MS-ADTS] describes both forms. For more information, see [MS-AUTHSOD] section 1.1.1.5.2, **Lightweight Directory Access Protocol (LDAP)** versions 2 and 3, **Kerberos**, and DNS.

Active Directory account creation mode: A type of account creation mode that retrieves and uses user accounts in a specific Active Directory Domain Services (AD DS) organizational unit.

anonymous access: A mechanism that does not require users to specify a user name or password for **authentication**.

anonymous authentication: An authentication mode in which neither party verifies the identity of the other party.

anonymous user: A user who presents no credentials when identifying himself or herself. The process for determining an anonymous user can differ based on the authentication protocol, and the documentation for the relevant authentication protocol should be consulted.

attachment: An external file that is included with an Internet message or associated with an item in a SharePoint list.

authentication: The ability of one entity to determine the identity of another entity.

authorization: The secure computation of roles and accesses granted to an identity.

back-end database server: A server that hosts data, configuration settings, and stored procedures that are associated with one or more applications.

basic authentication scheme: An HTTP-based **authentication** method that enables a protocol client to authenticate itself by passing a user identifier and password, as described in [RFC2617].

binary large object (BLOB): A discrete packet of data that is stored in a database and is treated as a sequence of uninterpreted bytes.

Central Administration site: A SharePoint site that an administrator can use to manage all of the sites and servers in a server farm that is running SharePoint Products and Technologies.

collection: A grouping of one or more EDM types that are type compatible.

column: See **field**.

configuration database: A database that is stored on a **back-end database server** and contains both persisted objects and site collection metadata for lookup purposes.

connection string: A series of arguments, delimited by a semicolon, that defines the location of a database and how to connect to it.

content database: A database that is stored on a **back-end database server** and contains stored procedures, site collections, and the contents of those site collections.

deployment: A collection of protocol clients and protocol servers (2) that belong to the same enterprise.

digital certificate: See the "digital certificate definition standard," as described in [X509].

directory service (DS): A service that stores and organizes information about a computer network's users and network shares, and that allows network administrators to manage users' access to the shares. See also **Active Directory**.

display name: A text string that is used to identify a principal or other object in the user interface. Also referred to as title.

document: An object in a **content database** such as a file, folder, **list**, or **site**. Each object is identified by a URI.

document library: A type of list that is a container for documents and folders.

domain: A set of users and computers sharing a common namespace and management infrastructure. At least one computer member of the set must act as a **domain controller (DC)** and host a member list that identifies all members of the domain, as well as optionally hosting the **Active Directory** service. The domain controller provides **authentication** of members, creating a unit of trust for its members. Each domain has an identifier that is shared among its members. For more information, see [MS-AUTHSOD] section 1.1.1.5 and [MS-ADTS].

domain controller (DC): The service, running on a server, that implements **Active Directory**, or the server hosting this service. The service hosts the data store for objects and interoperates with other **DCs** to ensure that a local change to an object replicates correctly across all **DCs**. When **Active Directory** is operating as Active Directory Domain Services (AD DS), the **DC** contains full NC replicas of the configuration naming context (config NC), schema naming context (schema NC), and one of the domain NCs in its **forest**. If the AD DS **DC** is a global catalog server (GC server), it contains partial NC replicas of the remaining domain NCs in its **forest**. For more information, see [MS-AUTHSOD] section 1.1.1.5.2 and [MS-ADTS]. When **Active Directory** is operating as Active Directory Lightweight Directory Services (AD LDS), several AD LDS **DCs** can run on one server. When **Active Directory** is operating as AD DS, only one AD DS **DC** can run on one server. However, several AD LDS **DCs** can coexist with one AD DS **DC** on one server. The AD LDS **DC** contains full NC replicas of the config NC and the schema NC in its **forest**. The domain controller is the server side of Authentication Protocol Domain Support [MS-APDS].

domain group: A container for security and distribution groups. A domain group can also contain other domain groups.

domain user: A user with an account in the domain's user account database.

email address: A string that identifies a user and enables the user to receive Internet messages.

end-user client: A computer on which an individual user is requesting specific file operations.

event receiver: A structured modular component that enables built-in or user-defined managed code classes to act upon objects, such as list items, **lists**, or content types, when specific triggering actions occur.

farm: A group of computers that work together as a single system to help ensure that applications and resources are available. Also referred to as server farm.

feature identifier: A GUID that identifies a feature.

field: A container for metadata within a SharePoint list and associated list items.

file: A single, discrete unit of content.

file system: A system that enables applications to store and retrieve **files** on storage devices. Files are placed in a hierarchical structure. The file system specifies naming conventions for files and the format for specifying the path to a file in the tree structure. Each file system consists of one or more drivers and DLLs that define the data formats and features of the file system. File systems can exist on the following storage devices: diskettes, hard disks, jukeboxes, removable optical disks, and tape backup units.

folder: A **file system** construct. File systems organize a volume's data by providing a hierarchy of objects, which are referred to as folders or directories, that contain files and can also contain other folders.

forest: In the **Active Directory** directory service, a **forest** is a set of naming contexts (NCs) consisting of one schema NC, one config NC, and one or more domain NCs. Because a set of NCs can be arranged into a tree structure, a **forest** is also a set of one or several trees of NCs.

forms authentication: An **authentication** method in which protocol clients redirect unauthenticated requests to an HTML form by using **HTTP**. If the protocol client authenticates the request, the system issues a cookie that stores the credentials or a key for reacquiring the identity. In subsequent requests, the cookie is submitted in request headers and the requests are authenticated and authorized by an ASP.NET event handler that uses the validation method that is specified by the protocol client.

front-end web server: A server that hosts webpages, performs processing tasks, and accepts requests from protocol clients and sends them to the appropriate back-end server for further processing.

group: A named collection of users who share similar access permissions or roles.

group object: A database object that represents a collection of user and group objects and has a **security identifier (SID)** value.

hierarchy: A logical tree structure that organizes the members of a dimension such that each member has one parent member and zero or more child members.

Hypertext Transfer Protocol (HTTP): An application-level protocol for distributed, collaborative, hypermedia information systems (text, graphic images, sound, video, and other multimedia files) on the World Wide Web.

Integrated Windows authentication: A configuration setting that enables negotiation of **authentication** protocols in Internet Information Services (IIS). Integrated Windows authentication is more secure than Basic authentication, because the user name and password are hashed instead of plaintext.

Kerberos: An **authentication** system that enables two parties to exchange private information across an otherwise open network by assigning a unique key (called a ticket) to each user that logs on to the network and then embedding these tickets into messages sent by the users. For more information, see [MS-KILE].

Lightweight Directory Access Protocol (LDAP): The primary access protocol for **Active Directory**. Lightweight Directory Access Protocol (LDAP) is an industry-standard protocol, established by the Internet Engineering Task Force (IETF), which allows users to query and update information in a **directory service (DS)**, as described in [MS-ADTS]. The Lightweight Directory Access Protocol can be either version 2 [RFC1777] or version 3 [RFC3377].

list: A container within a SharePoint site that stores list items. A list has a customizable schema that is composed of one or more fields.

list item: An individual entry within a SharePoint list. Each list item has a schema that maps to fields in the list that contains the item, depending on the content type of the item.

list view: A named collection of settings for querying and displaying items in a SharePoint list. There are two types of views: Personal, which can be used only by the user who created the view; and Public, which can be used by all users who have permission to access to the site.

login name: A string that is used to identify a user or entity to an operating system, directory service, or distributed system. For example, in Windows-integrated authentication, a login name uses the form "DOMAIN\username".

major version: An iteration of a software component, document, or list item that is ready for a larger group to see, or has changed significantly from the previous major version. For an item on a SharePoint site, the minor version is always "0" (zero) for a major version.

member: See OLAP member.

membership: The state or status of being a member of a member group. A membership contains additional metadata such as the privacy level that is associated with the membership.

NT LAN Manager (NTLM) Authentication Protocol: A protocol using a challenge-response mechanism for **authentication** in which clients are able to verify their identities without sending a password to the server. It consists of three messages, commonly referred to as Type 1 (negotiation), Type 2 (challenge) and Type 3 (authentication). For more information, see [MS-NLMP].

parent site: The site that is above the current site in the hierarchy of the site collection.

path component: Data that identifies a resource within the scope of a scheme and authority in a URI, as described in [RFC3986].

permission: A rule that is associated with an object and that regulates which users can gain access to the object and in what manner. See also **rights**.

permission level: A set of permissions that can be granted to principals or SharePoint groups on an entity such as a site, list, folder, item, or document.

placeholder: A character or symbol that is used in place of an actual value, text, or object. The actual value that the placeholder represents is unknown or unavailable at the current time, or is not displayed for security reasons.

pluggable security authentication: The ability to support alternate mechanisms for determining the identity of another entity.

query: A formalized instruction to a data source to either extract data or perform a specified action. A query can be in the form of a query expression, a method-based query, or a

combination of the two. The data source can be in different forms, such as a relational database, XML document, or in-memory object. See also **search query**.

result set: A list of records that results from running a stored procedure or query, or applying a filter. The structure and content of the data in a result set varies according to the implementation.

return code: A code that is used to report the outcome of a procedure or to influence subsequent events when a routine or process terminates (returns) and passes control of the system to another routine. For example, a return code can indicate whether an operation was successful.

rights: Tasks that a user is permitted to perform on a computer, site, domain, or other system resource. See also **permission**.

role: A symbolic name that defines a class of users for a set of components. A role defines which users can call interfaces on a component.

role definition: A named set of permissions for a SharePoint site. See also **permission level**.

round-robin load balancer: A resource management procedure where each process is assigned an equal portion of computer resources in a circular order.

scale-out: A method of adding computing resources by adding additional computers to the system, rather than increasing the computing resources on the computers in the system.

search provider: A component or application that provides data in response to a query. See also result provider.

search query: A complete set of conditions that are used to generate search results, including query text, sort order, and ranking parameters.

securable object: An object that can have unique security permissions associated with it.

security group: A named group of principals on a SharePoint site.

security group identifier: An integer that is used to uniquely identify a security group, distinguishing it from all other **security principals** and site groups within the same site collection.

security identifier (SID): An identifier for **security principals** in Windows that is used to identify an account or a group. Conceptually, the **SID** is composed of an account authority portion (typically a **domain**) and a smaller integer representing an identity relative to the account authority, termed the relative identifier (RID). The **SID** format is specified in [MS-DTYP] section 2.4.2; a string representation of **SIDs** is specified in [MS-DTYP] section 2.4.2 and [MS-AZOD] section 1.1.1.2.

security principal: A unique entity identifiable through cryptographic means by at least one key. A **security principal** often corresponds to a human user but can also be a service offering a resource to other **security principals**. Sometimes referred to simply as a "principal".

security scope: A tree structure of objects in which every object has the same security settings as the root.

server-relative URL: A relative URL that does not specify a scheme or host, and assumes a base URI of the root of the host, as described in [RFC3986].

Session Initiation Protocol (SIP): An application-layer control (signaling) protocol for creating, modifying, and terminating sessions with one or more participants. **SIP** is defined in [RFC3261].

Simple Mail Transfer Protocol (SMTP): A member of the TCP/IP suite of protocols that is used to transport Internet messages, as described in [RFC5321].

site: A group of related pages and data within a SharePoint site collection. The structure and content of a site is based on a site definition. Also referred to as SharePoint site and web site.

site collection: A set of **websites** that are in the same **content database**, have the same owner, and share administration settings. A site collection can be identified by a GUID or the **URL** of the top-level site for the site collection. Each site collection contains a top-level site, can contain one or more subsites, and can have a shared navigational structure.

site collection administrator: A user who has administrative permissions for a site collection.

SQL authentication: One of two mechanisms for validating attempts to connect to instances of SQL Server. In SQL authentication, users specify a SQL Server login name and password when they connect. The SQL Server instance ensures that the login name and password combination are valid before permitting the connection to succeed.

stored procedure: A precompiled collection of SQL statements and, optionally, control-of-flow statements that are stored under a name and processed as a unit. They are stored in a SQL database and can be run with one call from an application. Stored procedures return an integer return code and can additionally return one or more result sets. Also referred to as sproc.

store-relative URL: A URL that consists only of a path segment and does not include the leading and trailing slash.

thicket: A means of storing a complex HTML document with its related files. It consists of a thicket main file and a hidden thicket folder that contains a thicket manifest and a set of thicket supporting files that, together, store the referenced content of the document.

thumbnail: A miniature version of an image that is typically used to browse multiple images quickly.

topology: The structure of the connections between members.

transaction: The process of opening or creating an object on a server, and the subsequent committing of changes to the object by calling the required save function, at which time all changes to that instance of the object are either saved to the server, or discarded if a failure occurs before saving is finished successfully. Until successfully saved, changes are invisible to any other instances of the object.

Transact-Structured Query Language (T-SQL): A language that contains the commands that are used to manage instances of Microsoft SQL Server, create and manage all objects in an instance of SQL Server, and to insert, retrieve, modify, and delete all data in SQL Server tables. Transact-SQL is an extension of the language that is defined in the SQL standards that are published by the International Standards Organization (ISO) and the American National Standards Institute (ANSI).

trusted subsystem: A method of communication in which two-way trust is established between two server features. Each server feature communicates with the other feature by using an account that is authorized to perform privileged actions, such as retrieving files and settings.

Uniform Resource Locator (URL): A string of characters in a standardized format that identifies a document or resource on the World Wide Web. The format is as specified in [RFC1738].

user identifier: An integer that uniquely identifies a **security principal** as distinct from all other **security principals** and site groups within the same site collection.

user name: A unique name that identifies a specific user account. The user name of an account is unique among the other group names and user names within its own domain or workgroup.

user object: An object of class user. A user object is a security principal object; the principal is a person or service entity running on the computer. The shared secret allows the person or service entity to authenticate itself, as described in ([MS-AUTHSOD] section 1.1.1.1).

web application: (1) A container in a configuration database that stores administrative settings and entry-point **URLs** for **site collections**.

(2) A software application that uses **HTTP** as its core communication protocol and delivers information to the user by using web-based languages such as HTML and **XML**.

web application identifier: A GUID that identifies a web application.

Web Distributed Authoring and Versioning Protocol (WebDAV): The Web Distributed Authoring and Versioning Protocol, as described in [RFC2518] or [RFC4918].

Web Part: A reusable component that contains or generates web-based content such as **XML**, HTML, and scripting code. It has a standard property schema and displays that content in a cohesive unit on a webpage. See also Web Parts Page.

web server: A server computer that hosts websites and responds to requests from applications.

website: A group of related pages and data within a SharePoint site collection. The structure and content of a site is based on a site definition. Also referred to as SharePoint site and site.

XML: The Extensible Markup Language, as described in [XML1.0].

1.2 References

Links to a document in the Microsoft Open Specifications library point to the correct section in the most recently published version of the referenced document. However, because individual documents in the library are not updated at the same time, the section numbers in the documents may not match. You can confirm the correct section numbering by checking the Errata.

[MS-ADTS] Microsoft Corporation, "Active Directory Technical Specification".

[MS-FPSE] Microsoft Corporation, "FrontPage Server Extensions Remote Protocol".

[MS-SPO] Microsoft Corporation, "SharePoint Protocols Overview".

[MS-SYS] Microsoft Corporation, "Windows System Overview".

[MS-TDS] Microsoft Corporation, "Tabular Data Stream Protocol".

[MS-WDVSE] Microsoft Corporation, "Web Distributed Authoring and Versioning (WebDAV) Protocol: Server Extensions".

[MS-WDV] Microsoft Corporation, "Web Distributed Authoring and Versioning (WebDAV) Protocol: Client Extensions".

[MS-WSSFO2] Microsoft Corporation, "Windows SharePoint Services (WSS): File Operations Database Communications Version 2 Protocol".

[MS-WSSFOB] Microsoft Corporation, "Windows SharePoint Services (WSS): File Operations Database Communications Base Protocol".

[MS-WSSFO] Microsoft Corporation, "Windows SharePoint Services (WSS): File Operations Database Communications Protocol".

[MSDN-SHPTSDK4] Microsoft Corporation, "Microsoft SharePoint 2010 SDK", http://msdn.microsoft.com/en-us/library/office/ee557253(v=office.14).aspx

[MSDN-SHPTSDK] Microsoft Corporation, "Windows SharePoint Services 3.0 SDK", December 2007, http://msdn.microsoft.com/en-us/library/ms441339.aspx

[MSDN-SQLRBS] Microsoft Corporation, "Remote BLOB Store Provider Library Implementation Specification", Microsoft SQL Server 2008, http://msdn.microsoft.com/en-us/library/cc905212.aspx

[MSDN-WSSEBS] Microsoft Corporation, "External Storage of Binary Large Objects (BLOBs) in Windows SharePoint Services", SharePoint Services 3.0 SDK, http://msdn.microsoft.com/en-us/library/bb802976.aspx

[RFC2396] Berners-Lee, T., Fielding, R., and Masinter, L., "Uniform Resource Identifiers (URI): Generic Syntax", RFC 2396, August 1998, http://www.rfc-editor.org/rfc/rfc2396.txt

[RFC2518] Goland, Y., Whitehead, E., Faizi, A., et al., "HTTP Extensions for Distributed Authoring - WebDAV", RFC 2518, February 1999, http://www.ietf.org/rfc/rfc2518.txt

2 Functional Architecture

Windows SharePoint Services provides team-oriented collaboration websites , a platform for building web-based applications that use the Windows SharePoint Services collaboration features, and a framework for deploying and managing these sites and applications.

This section describes the architecture for delivering and supporting the framework in terms of computers, databases, external services, and protocols, and the architecture for supporting the collaboration features in terms of storage concepts within the framework.

A detailed discussion of security concepts is provided in section 2.9.

2.1 Overview

Windows SharePoint Services is designed to support a broad range of deployments. At a high level, four different types of systems are involved:

- The end-user client is a computer on which an individual user is requesting specific **file** and user operations. These requests are communicated using Hypertext Transfer Protocol (HTTP)-based protocols, including HTTP, WebDAV, and Microsoft FrontPage Server Extensions.
- The front-end web server is a computer that receives requests from an end-user client and provides Windows SharePoint Services capabilities by manipulating information stored in a database. These database requests are expressed in the T-SQL language and communicated using the Tabular Data Stream (TDS) protocol, as specified in [MS-TDS].
- The back-end database server is a computer that runs Microsoft SQL Server and responds to requests from the front-end web server.
- An **Active Directory** server responds to authentication requests from the end-user client, front-end web server, and back-end database server. These components could use an alternate authentication mechanism besides Active Directory.

These systems are shown in the following diagram.

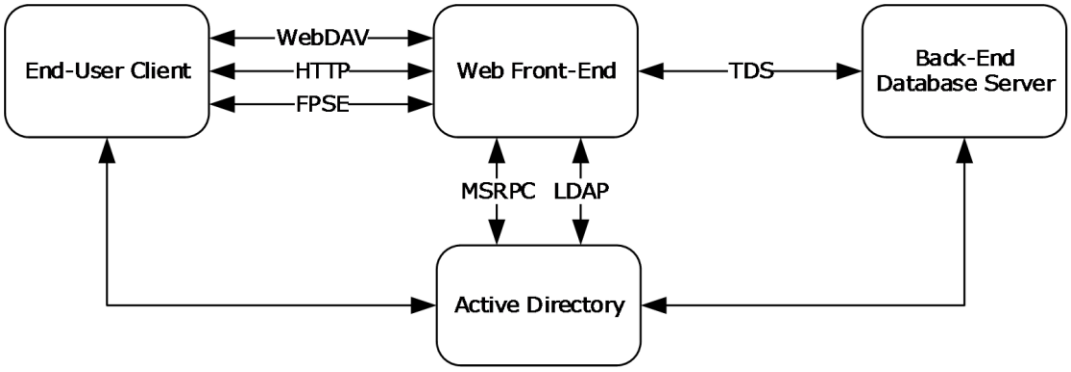

Figure 1: Intersystem protocol relationships

The front-end web server and back-end database server are considered part of the Windows SharePoint Services deployment; the end-user client is normally a user's desktop or laptop computer that connects to Windows SharePoint Services, and Active Directory is part of the generally available infrastructure. Both front-end web server and back-end database server can run on a single computer, so the simplest Windows SharePoint Services deployment could be just a single server. Alternatively, multiple instances of front-end web server and back-end database server computers can be installed for greater throughput and redundancy.

The following sections describe the architectural concepts pertaining to how Windows SharePoint Services uses the protocols specified by the member protocol specifications identified in section 2.2. These sections describe:

- How Windows SharePoint Services deployments can be scaled out with multiple computers, called a SharePoint **farm**. This **scale-out** is transparent to individual front-end web server computers as they respond to individual requests from an end-user client.
- The Windows SharePoint Services storage model, which allows for a variety of data management and organization techniques:
  - Storage for non-**file system** objects, such as **sites**, **site collections**, **lists**, and so on.
  - Storage for file system objects, such as files and **folders**.
  - Advanced storage concepts, such as **attachments**, **thickets**, ghosting, and so on.
- The SQL databases required for the operation of a Windows SharePoint Services deployment.

2.1.1 Scale-out Technologies

When a Windows SharePoint Services **deployment** is scaled out across multiple servers in a farm deployment, it uses two main technologies to increase throughput and availability.

1. **Network load balancing of the** front-end web server**:** Windows SharePoint Services supports network load-balancing technologies that distribute client requests across multiple servers in a farm. These individual front-end web servers are stateless; any front-end web server in the farm is prepared to handle any client request in the same way as any other front-end web server in the farm. By eliminating state or session information about the front-end web server, overall operational throughput can be increased simply by adding more front-end web servers. This stateless operation also makes the end-user experience more robust, because if a front-end web server fails, another front-end web server in the farm can handle future requests from the user. Front-end web servers do not communicate with one another in responding to client requests, but independently handle requests directed to them by the load-balancing technology.
2. **Vertical data partitioning across SQL databases:** As the deployment grows and the capacity of an individual server running SQL Server is fully consumed, additional back-end SQL resources can be deployed by adding additional servers that host completely separate **content databases**. Different site collections can be deployed into those separate content databases, and when a client request comes to a particular front-end web server, that front-end web server will fetch the site content strictly from the appropriate back-end database. This provides the ability to load-balance across multiple back-end resources, but does require manual placement of high-load sites into separate content databases.

The network load-balancing technology creates the need for each front-end web server in the server farm to behave identically to all other front-end web servers. The load balancer can be Windows Network Load Balancer, any one of a number of third-party hardware products, or even a simple custom **round-robin load balancer**.

2.1.2 Storage Architecture

Windows SharePoint Services provides a flexible model for storage which allows for a robust variety of data management and organization techniques. The following diagram presents a high-level view of the containers in this hierarchy. These containers provide important organizational and management tools for content stored in Windows SharePoint Services, and also form the core **securable object**.

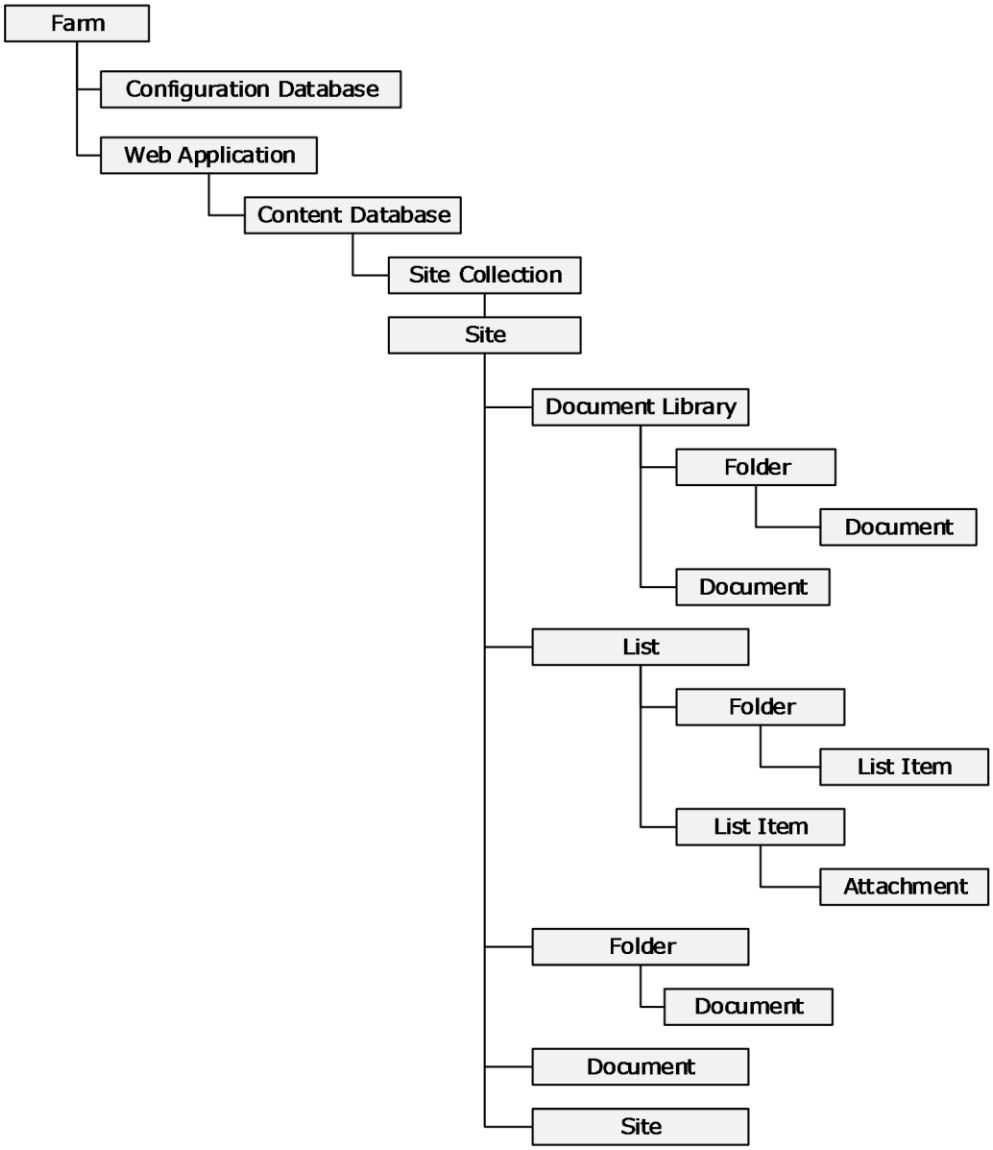

Figure 2: SharePoint storage object hierarchy

Each level of this hierarchy provides a specific set of management and deployment capabilities, and each item shown in the diagram is explained in detail in a later section.

2.1.2.1 Non-File System Objects

Multiple objects within Windows SharePoint Services are not directly tied to the representation of a file system. These objects are described in detail in later sections.

2.1.2.1.1 Farm

The SharePoint farm is the root container for an individual SharePoint deployment. This corresponds to a single **configuration database** (section 2.1.2.6) that describes the **topology** of the farm and global settings. Every object and container within a farm has a unique HTTP **Uniform Resource Locator (URL)**, which can be used to directly reference that object.

2.1.2.1.2 Web Application

The **web application (1)** is the container for grouping site collections within the context of the **web server**. A web application (1) is associated with one or more content databases that hold the content for site collections.

2.1.2.1.3 Site Collection

A site collection is a website within Windows SharePoint Services that can be created and managed directly by end users or, alternatively, more directly controlled by IT administrators. Characteristics of site collections include:

- They are the basic unit of scale for Windows SharePoint Services.
- They can be transparently organized.
- Multiple site collections can be distributed across multiple content databases.
- An individual site collection can only reside in one content database.

Site collections also provide a boundary for defining site groups (section 2.9.1.5) that can be given access to specific resources inside Windows SharePoint Services.

2.1.2.1.4 Site

A site is a container within a site collection that allows for delegated administration where certain actions, such as managing **permissions** to content within the site, can be delegated by the **site collection administrator** to the site administrator, who will have more direct knowledge of the business needs and allow for more efficient administration. A site collection will only have one site at the root, but can have many nested sites under the root site.

2.1.2.1.5 List

A list is a location within a site that maintains a configurable data structure where end-user data can be stored. A list is composed of multiple columns that define the structure for the data stored within the list.

2.1.2.1.6 List Item

A **list item** is a unique row within a list where individual end-user data elements are stored. A list item follows the structure of **columns** as defined by the list that contains the list item.

2.1.2.2 File System Objects

An important part of Windows SharePoint Services is the file system abstraction that it presents over data. Windows SharePoint Services uses basic file system concepts, such as files and folders. The following sections describe the Windows SharePoint Services objects that are exposed as file system concepts.

2.1.2.2.1 Document Library

A **document library** is a type of list specifically designed to be a location within a site where end-user **documents** are stored.

2.1.2.2.2 Folder

A folder is an organizational tool used within a site or document library that enables easier browsing and navigation. Within a document library, folders appear similar to their equivalent type within a file system.

2.1.2.2.3 Document

A document is an individual file that a user might create, read, or update using an authoring application. Documents can be stored within a site, document library, folder, or as an attachment.

2.1.2.3 Advanced Storage Concepts

In addition to basic file system concepts, Windows SharePoint Services provides a variety of specialized file system objects and operations as described in the following sections.

2.1.2.3.1 Attachment

Many list structures within Windows SharePoint Services can be configured to allow attachments, which allows multiple files to be included with each list item.

2.1.2.3.2 Thickets

For some complex HTML documents, Windows SharePoint Services provides the capability to store the document separated into its component sibling documents. This group of documents is treated like a single document for most file operations, but is stored as a set of related documents within the file store location.

2.1.2.3.3 Ghosting

Sites and lists tend to share a relatively small set of common system files across many instances within the back-end database server. To avoid having to store these files repeatedly for every site or list in the back-end database server, Windows SharePoint Services allows files to be ghosted, meaning that the content of the file is not actually stored in the back-end database server. Instead, a reference to a location on the front-end web server where the source of the file can be found is stored in the file metadata.

2.1.2.3.4 Versioning

Windows SharePoint Services can be configured on a per-list or per-document library basis to store multiple versions of documents. If versioning is configured for a storage location, each new version of a document is stored with an incrementing version number that can be either in the form "Major Version Number" (#) or "Major Version Number.Minor Version Number" (#.#).

If Major Version Number.Minor Version Number version numbering is used, individual documents start their numbering at 0.1, and the version can be promoted to a **major version** using the Publishing feature. Prior versions of a document can be retrieved by users with the appropriate access rights.

2.1.2.3.5 Publishing

Windows SharePoint Services can be configured on a per-list or per-document library basis to allow publishing features with documents. If publishing features are configured for a storage location, each version of a document can be configured as a draft, and therefore, not be available to be viewed by users without the appropriate access rights for viewing unpublished versions.

2.1.2.3.6 Document Property Promotion

The columns within a document library can be populated with data directly extracted from the document. A document can contain properties (usually metadata about the document such as author

or title) which can be associated with columns on a one property to one column basis. When a document is uploaded, the metadata is extracted from the document, and these associated columns are populated with the data. Conversely, Windows SharePoint Services has the capability to demote the associated column data back into the document properties if changes have occurred within the document library.

2.1.2.3.7 Large File Access

When dealing with very large files, obtaining the complete file contents in a single buffer as part of one operation can be a burden on system resources. Instead, the back-end database server provides functionality to return files larger than a specified size to the front-end web server in a series of smaller chunks that can be processed more smoothly.

2.1.2.3.8 BLOB Storage Outside the Content Database

Windows SharePoint Services provides the ability to allow file metadata to be stored by Windows SharePoint Services in a content database, while allowing the actual file contents to be stored in an external file system. The term **binary large object (BLOB)** refers to the SQL Server concept of unstructured, binary data streams that are commonly associated with files.

Windows SharePoint Services does not natively provide a BLOB store. There are two APIs that allow a BLOB storage provider to be built and registered with Windows SharePoint Services:

1. **External BLOB Storage:** For more information about External BLOB Storage, see [MSDN-WSSEBS]. Windows SharePoint Services 3.0 and wss4 support External BLOB Storage providers.
2. **Remote BLOB Storage:** For more information about Remote BLOB Storage, see [MSDN-SQLRBS]. Microsoft SharePoint Foundation 2010 supports Remote BLOB Storage providers.

> **Note** In SharePoint Foundation 2010, the use of Remote BLOB Storage is recommended over External BLOB Storage.

2.1.2.4 SQL Databases

The content database and configuration database are two core types of databases that are required for the operation of a Windows SharePoint Services deployment. These databases are considered "internal" to Windows SharePoint Services, which does not support users, developers, or system administrators directly accessing or manipulating content in these databases. This is true regardless of whether the Windows SharePoint Services deployment is running on a single server or across multiple computers.

Instead, Windows SharePoint Services exposes a full set of APIs to manage access to this data. For more information regarding use of these APIs in Windows SharePoint Services 3.0, see [MSDN-SHPTSDK]. For more information regarding use of these APIs in Microsoft SharePoint Foundation 2010, see [MSDN-SHPTSDK4].

2.1.2.5 Content Databases

A content database stores and manages end-user content. Each web application (1) will have one or more content databases. The Windows SharePoint Services content database stores the data and documents that end users have associated with their SharePoint site. The contents of this database are fully normalized to efficiently perform the kinds of operations required for the high-scale, highly configurable system previously described. The role of a Windows SharePoint Services front-end web server is to fetch the appropriate data from these multiple locations within SQL using the appropriate set of **queries** and **stored procedures** and to correctly interpret and map the results into a correct response to the end-user client. The queries and stored procedures are specified in [MS-WSSFOB] for

Windows SharePoint Services 2.0, [MS-WSSFO] for Windows SharePoint Services 3.0, and [MS-WSSFO2] for SharePoint Foundation 2010.

2.1.2.6 Configuration Database

The configuration database describes the topology of the farm and global settings. Each farm will have only one configuration database. The configuration database is essentially the definition of the Windows SharePoint Services deployment, for either a single instance of Windows SharePoint Services or for a farm. In the farm context, the configuration database contains information representing the global settings that are required to provide consistent operation across all servers within the farm, and to map requests to particular content databases. The configuration database allows only restricted access, as described in section 2.9.2.2. In a typical setup, the configuration database contents can only be modified from the **Central Administration site**, while run-time web applications (1) have only read access to these configuration settings.

2.1.2.6.1 Site Map

In addition to a description of the global topology, the configuration database also stores a site map, which is a mapping of all site collections to the individual content databases that contain the end-user content for the site collections. The following diagram shows how those URLs can be mapped to individual content databases.

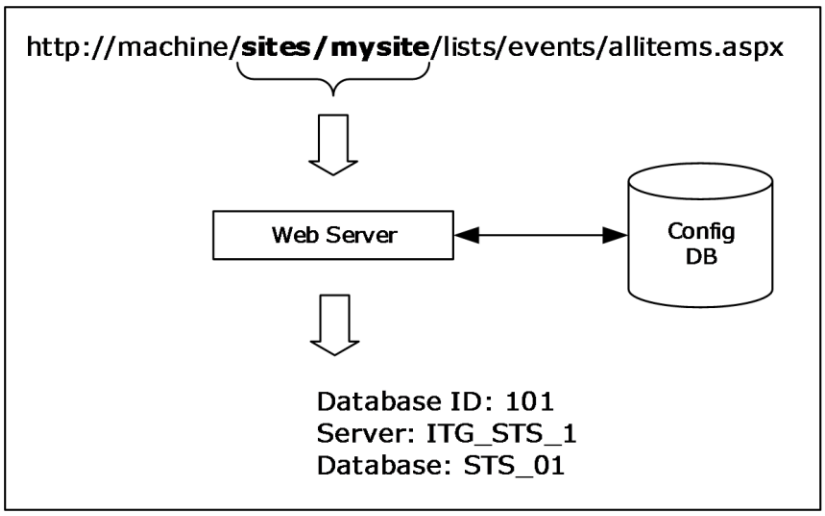

Figure 3: Determining the Site Collection URL

This mapping can use a section of the URL (in this case "sites/mysite") to map to an individual back-end database server and content database. This information is stored as a **server-relative URL**, so this mapping is effective across multiple, different names for the server. The server addresses http://machine, http://localhost, or http://157.55.234.184 are all equivalent.

The Site Collection Lookup (section 2.1.2.6.1.1) describes the process of looking up a site collection from the site map, and determining the **connection string** to the content database that holds the site collection's end-user generated content.

2.1.2.6.1.1 Site Collection Lookup

When a request is made for content at a specific URL, Windows SharePoint Services determines to which site collection the URL is referring, as well as the connection string of the content database holding the content for that site collection. This is accomplished in the following steps.

Step 1: Web Application Lookup

Site collection lookup begins by examining the portion of the incoming URL beginning with the Scheme Component and ending with the Authority Component (for example, http://example.com:80). Scheme and Authority are defined in [RFC2396] sections 3.1 and 3.2. This URL is compared against a stored set of web application (1) URLs. If one of the URLs in the list matches the incoming URL, the associated web application (1) is used for the remainder of the operation.

Step 2: Prefix Matching

Web applications (1) contain a set of site collection prefixes. These prefixes are URL **path components** that are used to determine which portion of the incoming URL path component is the server-relative URL of the site collection. This is done by matching all of the prefixes against the start of the path component of the incoming URL. If more than one prefix matches the beginning of the incoming URL path component, the longest matching prefix is used. A web application (1) can contain any combination of the two types of prefix:

- **Explicit Prefixes**: An explicit prefix indicates that the portion of the path component up to and including the prefix is included in the site collection server-relative URL. For example, if a user requests http://example.com/sitename/web/list/document.htm, and if the web application (1) corresponding to http://example.com contains an explicit prefix named "sitename", then "/sitename" is the server-relative URL of the site collection.

Incoming URL	Web application explicit prefixes	Resulting Site Collection Server-Relative URL
http://example.com/a/b/c.htm	"a"	"/a"
http://example.com/a/b/c.htm	"a", "a/b"	"/a/b"
http://example.com/a/b.htm	"a", "a/b"	"/a"
http://example.com/a/b.htm	"c"	<No Match>
http://example.com/a/b.htm	""	"/"

- **Wildcard Prefixes**: A wildcard prefix indicates that the portion of the path component up to and including the first path component segment following the prefix is included in the site collection name. For example, if a user makes a request for http://example.com/sites/sitename/web/list/document.htm, and if the web application (1) corresponding to http://example.com contains a wildcard prefix named "sites", then "/sites/sitename" is the server-relative URL of the site collection.

Incoming URL	Web application wildcard prefixes	Resulting Site Collection Server-Relative URL
http://example.com/a/b/c/d.htm	"a", "a/b"	"/a/b/c"
http://example.com/a/b.htm	"a", "a/b"	<No Match>
http://example.com/a/b.htm	""	"/a"

Step 3: Site Collection Identifier Lookup

Once the site collection URL is determined, it is passed to the configuration database, along with the **web application identifier**. A site collection identifier is returned along with the identifier of the content in which the site **collection** content is stored. If the specified combination of site collection URL and web application identifier cannot be found in the configuration database, the site collection does not exist.

Some site collections are identified not by the path component of the URL, but by the URL Component (For example, "example.com:80"). These are known as Host Header Site Collections. If the web

application (1) cannot be identified from the Scheme and Authority Components of the incoming URL, site collection lookup assumes that the incoming URL refers to a Host Header Site Collection.

In this case, the Authority Component of the incoming URL is passed to the configuration database, which returns the corresponding site collection identifier. The site collection identifier is then passed back to the configuration database, which returns the identifier of the content database in which the site collection content is stored. If the specified Authority Component cannot be found in the configuration database, the site collection does not exist.

Step 4: Content Database Connection String Lookup

Once the content database identifier is known, a lookup occurs to determine connection string information about the content database. The following steps occur to generate this connection string:

- The content database identifier is passed to the configuration database, which returns the content database name and the identifier of the database service that is hosting the content database. If SQL authentication is intended to be used when connecting to the content database, the connection **user name** and password are also returned at this time.
- The identifier of the database service is then passed to the configuration database, which returns the name of the database service and the identifier of the server on which the database service is running.
- The identifier of the server is passed to the configuration database, which returns the address of the server.
- Finally, the server address, database service name, content database name, and optionally, the content database user name and password are combined to build the content database connection string.

2.2 Protocol Summary

The following table provides a comprehensive list of the member protocols of the Windows SharePoint Services File, Print, and User/Group Administration system.

Protocol name	Description	Short name
Windows SharePoint Services (WSS): File Operations Database Communications Base Protocol	This protocol specifies the communication between the front-end web server and the back-end database server that is used to satisfy requests involving file access and the administration of users and groups within Windows SharePoint Services 2.0.	[MS-WSSFOB]
Windows SharePoint Services (WSS): File Operations Database Communications Protocol	This protocol specifies the communication between the front-end web server and the back-end database server used to satisfy requests involving file access and administration of users and groups within Windows SharePoint Services 3.0.	[MS-WSSFO]
Windows SharePoint Services (WSS): File Operations Database Communications Version 2 Protocol	This protocol specifies the communication between the front-end web server and the back-end database server used to satisfy requests involving file access and administration of users and groups within SharePoint Foundation 2010.	[MS-WSSFO2]
Web Distributed Authoring and Versioning (WebDAV) Protocol: Client Extensions	The client extensions in this protocol extend the WebDAV Protocol, as specified in [RFC2518], by introducing new headers that both enable the file types that are not currently manageable and optimize protocol interactions for file system clients. These WebDAV Protocol: Client Extensions do not introduce new functionality into the WebDAV Protocol, but instead optimize processing and eliminate the need for	[MS-WDV]

Protocol name	Description	Short name
	special-case processing.	
Web Distributed Authoring and Versioning (WebDAV) Protocol: Server Extensions	The server extensions in this protocol extend WebDAV by introducing new HTTP request and response headers that both enable the file types that are not currently manageable and optimize protocol interactions for file system clients. These extensions also introduce a new WebDAV method that is used to send **search queries** to disparate **search providers**.	[MS-WDVSE]
FrontPage Server Extensions Remote Protocol	This protocol specifies a set of server extensions that can be used to augment a basic HTTP server. These extensions provide file server functionality similar to WebDAV, allowing a website to be presented as a shared folder. The use of WebDAV is recommended over the FrontPage Server Extensions Remote Protocol. The SharePoint Team Services dialogview is an application of the FrontPage Server Extensions Remote Protocol that is addressed in the FrontPage Server Extensions Remote Protocol Specification [MS-FPSE] because it has certain behaviors apart from the normal FrontPage Server Extensions Remote Protocol communications. The purpose of the dialogview is to allow a client to display a server-rendered HTML-based rendering of the files located on a particular website.	[MS-FPSE]

2.3 Environment

The following sections identify the context in which the system exists. This includes the systems that use the interfaces provided by this system of protocols, other systems that depend on this system, and, as appropriate, how components of the system communicate.

2.3.1 Dependencies on this System

None of the systems in Windows Server 2003 operating system, Windows Server 2008 operating system with Service Pack 2 (SP2), or Windows Server 2008 R2 operating system that are used to deliver file, print, user administration, or group administration services depend on this system.

2.3.2 Dependencies on Other Systems/Components

The Windows SharePoint Services File, Print, and User/Group Administration system depends on the following systems:

- Windows System: [MS-SYS]
- Tabular Data Stream Protocol: [MS-TDS]
- Active Directory: [MS-ADTS]

Windows SharePoint Services 3.0 depends on the following components to function:

- Windows Server 2003 with Service Pack 1 (SP1), Windows Server 2003 with Service Pack 2 (SP2), Windows Server 2003 with Service Pack 3 (SP3), and Windows Server 2003 R2 operating system
  - Internet Information Services (IIS) 6.0
- Microsoft .NET Framework 3.0 or Microsoft .NET Framework 3.5

- Microsoft ASP.NET 2.0

SharePoint Foundation 2010 depends on the following systems/components to function:

- Windows Server 2008 and Windows Server 2008 R2
  - Internet Information Services (IIS) 7.0
- .NET Framework 3.5
  - ASP.NET 2.0
- Microsoft Forefront Unified Access Gateway
- Microsoft SQL Server 2008 Express Edition with Service Pack 1

2.3.2.1 Domain Controller/Directory Service

In addition, Windows SharePoint Services can communicate with an Active Directory **domain controller (DC)** to provide **authentication** services that enable the User/Group Administration functions described in this document. This domain controller provides an **LDAP**-enabled **directory service (DS)** that stores user information, such as name and **email address**.

2.4 Assumptions and Preconditions

This section briefly documents the assumptions and preconditions required by the system. The scope of this discussion is intended to be implementation-independent and is limited to the system level of Windows SharePoint Services.

- The Windows SharePoint Services server(s) is reachable by external clients via an established IP address (or IP addresses).
- The Windows SharePoint Services server(s) functional components are started collectively and the Windows SharePoint Services server(s) accepts client requests.
- The Windows SharePoint Services front-end web servers can reach back-end database servers and have appropriate permissions to access data in the content database and the configuration database.
- The Windows SharePoint Services front-end web server and back-end database server are matching versions, or within an acceptable range of versions. For more information about versioning, see section 2.6.
- In the case where Active Directory is used to provide authentication, the directory service (DS) is accessible to the Windows SharePoint Services server. Any intermediate firewalls, routers, or connection points between components of the system have all required ports and gateways open for communication between them.

2.5 Use Cases

The following use case is provided only for an understanding of the Windows SharePoint Services File, Print, and User/Group Administration system. It is not intended to be a thorough and complete modeling of the system for implementation purposes.

2.5.1 Creating a SharePoint Document Library File from the Client Console

This use case describes the simplest way to create a file using the protocols covered in this system. The actor in this case is the user who is creating the text file hello.txt in a Windows SharePoint

Services document library. The text file contains the text "hello". For details regarding a scenario of this type, see the example Create File from Client in section 3.4.

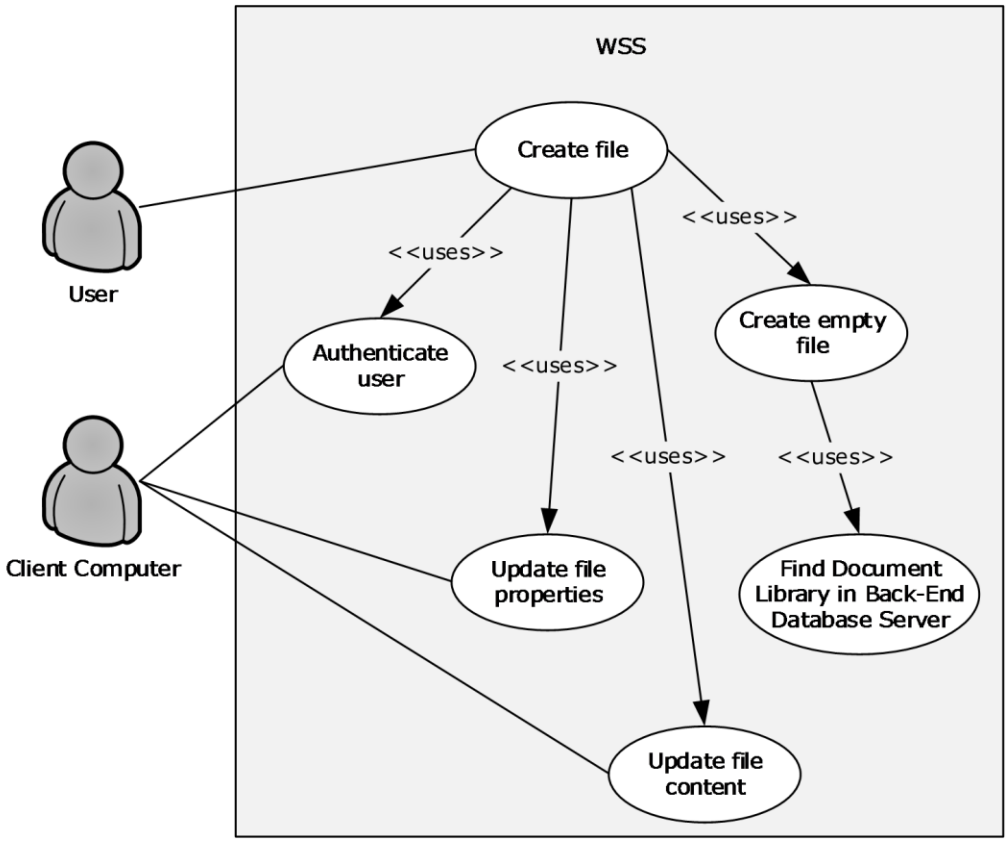

Figure 4: Create file from client computer

Preconditions

- The user has read/write access permissions to an existing SharePoint document library called http://server/site/doclib.
- The user is logged on to a client computer running the Windows 7 operating system<1> with an authenticated Windows session, and can access the Windows SharePoint Services site containing the document library.
- From a command prompt window, the user types the following command:

```
echo hello >\\server\site\doclib\hello.txt
```

Main Flow

1. The user types the `echo` command and presses **Enter**.
2. The Windows SharePoint Services front-end web server authenticates the user.
3. Windows SharePoint Services finds the location of the document library and verifies that the user has the appropriate access permissions.
4. Windows SharePoint Services creates an empty file in the document library and confirms successful creation of the file back to the client.

5. The client updates the file properties and the file contents.

Error Scenarios

- The user does not have the appropriate access permissions: The client notifies the user that access is denied.
- The client cannot connect to the Windows SharePoint Services front-end web server: The client notifies the user of the connection error.
- The client cannot update the file properties or file contents: The client notifies the user of the file write properties or file contents write error.

2.6 Versioning, Capability Negotiation, and Extensibility

The Windows SharePoint Services front-end web server and Windows SharePoint Services back-end database server perform explicit version verifications in Windows SharePoint Services 3.0 and SharePoint Foundation 2010.

The client calls the proc_GetVersion stored procedure to retrieve version information from the server and to decide whether it ought to connect to the database. The proc_GetVersion stored procedure is described in [MS-WSSFO] section 3.1.5.39 for Windows SharePoint Services 3.0, and in [MS-WSSFO2] section 3.1.5.44 for SharePoint Foundation 2010.

The version information is stored in the Versions table, which is described in [MS-WSSFO] section 2.2.7.11 for Windows SharePoint Services 3.0, and [MS-WSSFO2] section 2.2.7.11 for SharePoint Foundation 2010.

Windows SharePoint Services front-end web server initializes the connection to the back-end database server according to the following steps:

1. When the front-end web server was added to the SharePoint farm, the administrator gave a connection string to the configuration database.
2. In SharePoint Foundation 2010, the front-end web server verifies that the configuration database is the correct version by calling proc_GetVersion with each of the version identifiers specified in table B of [MS-WSSFO2] section 3.1.5.44. The front-end web server also ensures that the version numbers are within the acceptable range defined in that same table. If one or more versions are not within the acceptable range, the front-end web server disconnects from the configuration database.
    - In Windows SharePoint Services 3.0, the version of the configuration database is not checked; the connection initialization operation proceeds directly to the next step.
3. Once the front-end web server verifies that it is connected to a configuration database with an appropriate version, the front-end web server gets the connection information for the content database that is required to respond to the URL request being handled. For more details on the URL Site Map Lookup, see sections 2.1.2.6 and 2.1.2.6.1.
4. When the front-end web server has the connection string to the content database, the front-end web server connects to the content database and verifies that it is the correct version by calling proc_GetVersion with each of the version identifiers specified in table A of [MS-WSSFO] section 3.1.5.39 for Windows SharePoint Services 3.0, or in [MS-WSSFO2] section 3.1.5.44 for SharePoint Foundation 2010, and ensures that the version numbers are within the acceptable range defined in the same table. If one or more versions are not within the acceptable range, the front-end web server disconnects from the content database.
    - Windows SharePoint Services 3.0 uses only one of the version identifiers. For more information, see the product behavior note for proc_GetVersion in [MS-WSSFO2] section 3.1.5.44 regarding SharePoint Foundation 2010.

5. The validation result is cached in the front-end web server process. When the process restarts, the validation will be performed again.

The acceptable range of specified version numbers can change when Windows SharePoint Services is updated through a service pack or other release.

2.7 Error Handling

There are no system-level error handling behaviors. In general, for errors returned as part of a protocol in this system, the technical documents for those protocols describe what the error means for the system when they are defined. How these errors are handed, based on the protocol description, is left to the implementer.

2.8 Coherency Requirements

This system has no special coherency requirements.

2.9 Security

This section describes two core aspects of the Windows SharePoint Services security model: authentication and **authorization**.

Authentication is the part of the system that determines the current user's identity. This is the first step in managing the security of the system. Windows SharePoint Services uses the authentication mechanism from an underlying platform, such as Internet Information Services (IIS) and ASP.NET, to authenticate users.

Windows SharePoint Services supports all of the authentication modes that IIS and ASP.NET support, including Active Directory, **forms authentication**, and WebSSO authentication. In Active Directory authentication mode, IIS authenticates the user, using **basic authentication scheme**, **digital certificate**, **NT LAN Manager (NTLM) Authentication Protocol**, or **Kerberos**. In other authentication modes, Windows SharePoint Services relies on ASP.NET authentication modules to authenticate users, which can also be created by third-party developers, such as **FormsAuthenticationModule** or **ADFSAuthenticationModule**. For more information about Active Directory authentication, see section 2.9.2.

Authorization in Windows SharePoint Services identifies which permissions are granted to which users on a given object. When a web request (or some object model API code) attempts to access an object inside Windows SharePoint Services, and the caller has been authenticated, the authorization code is called to identify whether the access can be granted. In a **trusted subsystem** model, the front-end web server uses the IIS application identity account to access the contents in the content database, on behalf of the user, to access content rather than the account of the user who is using the site. For more information, see section 2.9.2.2. Therefore, the permissions check has to happen before Windows SharePoint Services returns any page content back to the user.

The following sections describe the basic concepts pertaining to authorization.

2.9.1 Authorization for User and Group Administration

After the user has been identified (authenticated), Windows SharePoint Services controls the user's authorization and determines which permissions are granted to that user on a given object.

Windows SharePoint Services supports a number of security-related operations to control access to content stored in Windows SharePoint Services. These operations are built around a few core concepts defined in Windows SharePoint Services, which are described in detail in the following sections.

2.9.1.1 Individual User Permissions (Rights)

Individual permissions, also known as **rights**, grant the ability to perform specific actions. For example, the View Items permission grants a user the ability to view items in a list. Windows SharePoint Services has a fixed set of permissions that can be granted to users. The full specification of the Windows SharePoint Services Rights Mask is provided in the Windows SharePoint Services Rights Mask section of [MS-WSSFOB] section 2.2.2.10 for Windows SharePoint Services 2.0, [MS-WSSFO] section 2.2.2.13 for Windows SharePoint Services 3.0, and in [MS-WSSFO2] section 2.2.2.14 for SharePoint Foundation 2010. The permissions directly related to file services scenarios are:

- Add Items
- Edit Items
- Delete Items
- View Items
- Open Items
- Browse Directories

In addition, the following set of permissions relate to permissions control:

- Manage Permissions
- Create Groups
- Enumerate Permissions
- Open (site, web, list, folder)

2.9.1.2 Permission Level (Role)

A **role** is a predefined set of permissions that grants users permission to perform related actions. Roles are defined at the site level, where a site can inherit roles from its **parent site** or have roles unique to it. All permissions in Windows SharePoint Services are managed through roles and all users will have roles. Rights are never directly assigned to a user. The default Windows SharePoint Services permission levels, or roles, are:

- Limited Access
- Read, Contribute
- Design
- Full Control

For example, the Limited Access role includes permissions that allow users to view specific lists, document libraries, list items, folders, or documents, when given the appropriate permissions.

It is also possible to add custom **role definitions** to the collection of roles, to include the specific set of rights required for the role, or to remove role definitions. For example, a specific scenario might require a user role where the user cannot see previous versions of a document. To achieve this, it is possible to create a custom contributor role where the View Versions and Delete Versions rights have been removed.

For more information about creating and removing roles, see [MSDN-SHPTSDK] for Windows SharePoint Services 3.0, and [MSDN-SHPTSDK4] for SharePoint Foundation 2010.

2.9.1.3 User

A user is an identity associated with a user account that can be authenticated to Windows SharePoint Services. **Permission levels** can be directly assigned to users. After a user has been authenticated, that user's identity is represented by a Windows SharePoint Services user token, and their permissions are represented by the roles to which they are assigned, as described in section 2.9.2. Role assignment is per site, where Windows SharePoint Services tracks which users (or groups, as described in section 2.9.1.4) are assigned to which roles for each site (2). A user's complete set of permissions is an aggregation of two sets of roles:

- The roles of which the user is a direct **member**.
- The roles that the user acquires by being a member of a **group** or site group (section 2.9.1.5).

2.9.1.4 Group

A group is an identity associated with a group of users within Active Directory (Windows **security group**). As far as account management is concerned, Windows SharePoint Services treats groups similarly to user accounts. When a user interacts with a Windows SharePoint Services environment, their Active Directory group **membership** is determined and their membership within a group is used to determine the effective role of the user.

2.9.1.5 Site Group

A Windows SharePoint Services site group is a named logical grouping of user or group accounts. A site group can be set to specific roles or have rights granted to it. Each Windows SharePoint Services site group is assigned a default role, but the role for any site group can be changed as necessary. Some predefined Windows SharePoint Services site groups are as follows:

- Site Owners
- Site Members
- Site Visitors

Windows SharePoint Services group (2) memberships are stored in SQL table named GroupMemberships. Each group is assigned an identifier that is unique within that site collection. The use of groups can enable easier security management. When a large number of users have to be assigned the same role, administrators can easily create a Windows SharePoint Services group (2) and assign those users as members (3) and simply grant permissions to the group rather than to each individual. Similarly, administrators can add new users to existing groups as a means of quickly giving users appropriate permissions. For more information about creating Windows SharePoint Services groups (2) and adding users to existing groups, see:

- [MSDN-SHPTSDK] for Windows SharePoint Services 3.0
- [MSDN-SHPTSDK4] for SharePoint Foundation 2010
- proc_SecCreateSiteGroup and proc_SecAddUserToSiteGroup in [MS-WSSFOB] for Windows SharePoint Services 2.0
- proc_SecCreateSiteGroup and proc_SecAddUserToSiteGroup in [MS-WSSFO] for Windows SharePoint Services 3.0
- proc_SecCreateSiteGroup and proc_SecAddUserToSiteGroup in [MS-WSSFO2] for SharePoint Foundation 2010

Windows SharePoint Services groups (2) cannot be nested inside of each other. However, a Windows SharePoint Services group (2) can contain Active Directory groups as members (3).

Windows SharePoint Services groups (2) are themselves a securable object in Windows SharePoint Services with specific permissions to manage them, as described in section 2.9.1.2.

2.9.1.6 Securable Object

Users are assigned a permission level for a specific securable object: a site, library, folder, or document. By default, permissions for a site, library, or document are inherited from the parent site or library.

Folders within lists are securable objects in that they derive their permissions from the underlying list items they contain. Pages that do not belong to any list are not securable objects; such pages always share the same permissions as their parent site. Attached files (attachments) and **thumbnail** files are also not securable objects; they always share the same permissions as their associated list item.

Each securable object gets its security permissions from its **access control list (ACL)** and other security metadata (for example owner info, checkout state, and so on). The security permissions can be unique, or can be inherited from the parent of the object.

2.9.1.7 Scope

A **security scope** represents a URL subtree in Windows SharePoint Services that shares the same permissions. A user creating an item in Windows SharePoint Services could choose to give the item its own specific permission requirements or specify that it can inherit permissions from its parent. If the item has its own permissions, the item and its descendants form a scope, with the created item being the root of that scope. If the item inherits permissions, the item belongs to a larger scope that also contains its parent. A scope cannot span more than one site collection; however, it can span multiple sites within a site collection.

2.9.1.8 Inheritance

As mentioned earlier, an item in Windows SharePoint Services can have its own specific permissions. However, default permissions for an item within a site are inherited from that site. This inheritance can be broken for any securable object at a lower level in the site **hierarchy** by creating a unique permissions assignment for that securable object.

For example, editing the permissions for a document library breaks its permission inheritance from its site. However, the inheritance is broken only for that particular document library; the rest of the permissions for the site remain unchanged. An object can be reverted to inheriting permissions from its parent list or site at any time.

Sites are themselves a securable object to which permissions can be assigned. Sites contained within other sites can be configured to inherit permissions from a parent site or to create unique permissions for that particular site. If a child site inherits permissions from its parent, that set of permissions is shared with the child site. This effectively means that owners of the sites that inherit permissions from a parent site can change the permissions of the parent. To allow control of the permissions for the child site alone, the child site stops inheriting permissions, restricting the owner of the child site to only making changes to the permissions of the child site.

Creating unique permissions for a site stops permission inheritance. The groups, users, and permission levels from the parent site are copied to the child site and then the inheritance is broken. Reverting a site back from unique permissions to inherited permissions causes users, groups, and permission levels to once again be inherited and removes any users, groups, or permission levels that were uniquely defined in the site while inheritance was broken.

Permission levels (roles) can also be inherited. By default, permissions are defined such that the Read permission level is the same regardless of the object to which it is applied. This type of inheritance can also be broken by editing the permission level. For example, an administrator might not require the Read permission level on a particular site (2) to include the Create Alerts permission.

For more information about inheritance, see:

- [MSDN-SHPTSDK] for Windows SharePoint Services 3.0
- [MSDN-SHPTSDK4] for SharePoint Foundation 2010
- proc_SecChangeToInheritedWeb and proc_SecChangeToUniqueWeb in [MS-WSSFOB] for Windows SharePoint Services 2.0
- proc_SecChangeToInheritedWeb and proc_SecChangeToUniqueScope in [MS-WSSFO] for Windows SharePoint Services 3.0
- proc_SecChangeToInheritedWeb and proc_SecChangeToUniqueScope in [MS-WSSFO2] for SharePoint Foundation 2010

2.9.1.9 Anonymous

Windows SharePoint Services allows a specific type of access where the user is not uniquely authenticated, and thus is unknown to Windows SharePoint Services. Such a user is referred to as the "Anonymous" user, or as having "Anonymous" access. The availability of **anonymous access** is controlled at the web application (1) level of Windows SharePoint Services. If anonymous access is allowed for the web application (2), then for example, site administrators can decide whether to:

- Grant anonymous access to a site
- Grant anonymous access only to lists and libraries
- Block anonymous access to a site (2) altogether

Anonymous access relies on the **anonymous user** account on the web server. This account is created and maintained by Microsoft Internet Information Services (IIS), not by Windows SharePoint Services. By default in IIS, the anonymous user account is IUSR_ComputerName. Enabling anonymous access essentially grants the anonymous account access to the Windows SharePoint Services site. Allowing access to a site, or to lists and libraries, grants the View Items permission to the anonymous user account. However, even with the View Items permission, there are restrictions to what anonymous users can do. For example, anonymous users cannot perform the following actions:

- Upload or edit documents into document libraries, including wiki libraries
- View the site (2) in My Network Places

When the user who is accessing Windows SharePoint Services items is anonymous, then the user identifier is null. In authentication (1) modes other than Active Directory, such as forms authentication, Windows SharePoint Services is also impersonating IUSR_ComputerName. In those cases, Windows SharePoint Services uses a **user identifier** generated from the ASP.NET Identity.Name value.

2.9.1.10 Anonymous Rights Mask (Anonymous Permissions Mask)

Use of an Anonymous Rights Mask (Anonymous Permissions Mask) is based on the concept of permissions for an anonymous user. The anonymous user can be given permissions via a role, or can be assigned direct permissions on an item in Windows SharePoint Services. Because, unlike other Windows SharePoint Services users, the anonymous user's permissions cross site collection boundaries and the user identifier is null (there is no "user" per se to have permissions), Windows SharePoint Services uses an anonymous permission mask at the security scope level to make a security decision regarding whether the anonymous user has access to items within that scope.

2.9.1.11 System Account

Some features in Windows SharePoint Services can trigger operations that require the ability to run with full permissions if the current user does not have permission to do so directly. For security reasons, the feature cannot reveal the identity of the account with such permissions. This concept is known as "run as system account". When an account runs with "run as system account" permissions, operations performed by this account are recorded as executed by the "system account". The login name for the system account is SHAREPOINT\system. For example, when a list item is created by the application pool identity account, it will show as "created by SHAREPOINT\system".

2.9.2 Authentication

Windows SharePoint Services supports **pluggable security authentication**, an extensibility mechanism provided by ASP.NET. By default, Windows SharePoint Services uses one of three authentication modes against a Windows **domain**:

- **Integrated Windows authentication**
- Basic authentication scheme
- **Anonymous authentication**

Specific deployments can use a custom authentication provider to authenticate end users against any third-party authentication system.

When used with Active Directory and a Windows domain, Windows SharePoint Services works with Active Directory for authentication of network accounts in the following contexts:

- Authentication of the requests from the end-user client. The front-end web server establishes a specific end-user identity for requests from the end-user client. The front-end web server evaluates that end-user identity against permissions associated with objects related to the request, to determine whether to execute the action for that request.
- Authentication of the Process Account from the front-end web server. The back-end database server establishes an identity for requests from the front-end web server. The back-end database server evaluates whether that identity has permissions to operate as a Windows SharePoint Services front-end web server for content stored in the back-end database server.
- Creating a site collection local user record for each logged-in user.
- Updating the site collection local user record to reflect a change in the user record in Active Directory.
- Selecting users and groups from the directory for the purposes of setting security access control lists (ACLs), as well as defining SharePoint groups (2).
- Creating Active Directory user accounts in **Active Directory account creation mode** to enable the creation of Active Directory accounts for Windows SharePoint Services users.

Section 2.9.2.1 specifies how Windows SharePoint Services uses the Active Directory Protocol [MS-ADTS] for the two types of authentication previously described.

2.9.2.1 Authentication of the Requests from the End-User Client

A typical scenario for authentication of the requests from the end-user client occurs when a user is logged in using an Active Directory **domain user** account that allows the user to access a network resource such as a Windows SharePoint Services site. When the user makes a page request to Windows SharePoint Services, IIS handles the request and authenticates the user. IIS could choose

one of the three methods: Integrated Windows authentication, basic authentication scheme, or anonymous authentication.

Once IIS has authenticated the user, IIS impersonates that user account for the thread handling the request. At this point, control is handed over to the Windows SharePoint Services code to fulfill the request from the end-user client. The request contains the unique address of a resource in Windows SharePoint Services (a page, a document, a list item, and so on), and its associated ACL. The ACL specifies which **security principal** has what permissions on this object. It is possible that the object inherits its permissions from an object higher in the container hierarchy (for example, see the figure in section 2.1.2).

In addition to the resource's address, the request contains the action that has to be performed on the object, such as Read, Write, Delete, Check-out, and so on. The Windows SharePoint Services authorization system performs the following steps to determine whether the requestor can perform the requested action on the request object.

1. When making the determination, the authorization function inspects the user's token (a data structure provided by Active Directory on the thread by IIS) containing the user and **security group identifier**. It compares this against the list of referenced SIDs in the site collection. As a performance optimization, once this comparison is made, only the SIDs in the user's token that are referenced in the site collection are used.

2. The authorization code then uses the truncated user token and compares it against each **access control entry (ACE)** in the requested object's ACL. An ACE contains the security principal and the action(s) which the principal can perform on that object. By matching all of the principals in the ACEs against the user's token, the list of actions that the user can perform on the requested object is determined.

3. The final step is to compare the requested action against the list of actions that the user can perform as determined by the authorization algorithm. If the requested action is present in the list of authorized actions for the user, the request is allowed; otherwise, the request is denied.

2.9.2.2 Authentication of the Process Account from the Front-End Web Server

When the front-end web server requires data from the back-end database server, it makes a request for that data via the appropriate protocol. The back-end database server has to validate whether the front-end web server has appropriate permissions to access the data in the back-end database server.

The Windows SharePoint Services worker process in the front-end web server runs under a service account identity defined in Active Directory. This account is assigned by the system or can be set by the administrator. The back-end database server validates the requestor's identity using either Integrated Windows authentication or SQL login. If Integrated Windows authentication is chosen, the request made to the back-end database server is done using the service account identity.

The back-end database server then uses the requestor's identity (from the SPUser object, not the requestor's account) to authorize the request.

2.9.2.3 Creating a Site Collection Local Record of the User

In addition to authorizing user requests, Windows SharePoint Services has to know the user's identity for other purposes, such as indicating who created a document, providing an email address to deliver an alert, or associating a picture with a discussion post.

To optimize back-end database server page rendering performance, some user information is copied from Active Directory and stored in the content database for that site collection. This allows the user's name, email address or **Session Initiation Protocol (SIP)** address, and picture to be used in rendering a page without making a call to Active Directory.

When the user makes a request to see the list of documents in a document library, the name of the user who last created or updated the file is displayed. To avoid calling Active Directory for each document just to get the appropriate user name, the user names are stored in a User Info table in the content database. This allows the back-end database server to determine the user name by doing a database Join instead of an off-computer remote call to Active Directory.

This User Info table is populated with an entry containing the user's SID, account name, name, email address, SIP address, title, and department. The entry is created when one of the following actions occur:

- A user is given access using the Site Administration pages.
- A user is referenced by a list item (for example, Task is assigned to "User A").
- A user uploads or creates a document.
- A user visits the site by requesting a document, page, or item from the site.

All of the fields for the user entry are found in the user's Active Directory **user object** record. If the Active Directory user object record is not populated with any information, the resulting entry in the Windows SharePoint Services User Table also contains nothing. This is a one-time occurrence; the table is not regularly updated against Active Directory.

2.9.2.4 Updating the Site Collection Local User Record (Account Migration)

During a user's life cycle, a number of user attributes can change in Active Directory. Some common changes are name (user has legally had their name changed), domain (the user was migrated from one Active Directory domain to another), or **forest** (user was migrated from one Active Directory forest to another).

Because there is no automatic way to synchronize the user information in the User Info table of the content database with the Active Directory, a MigrateUser command line option is provided. The most common use of this command is for updating the user record when the user has been migrated from one domain to another within the same forest, or from one forest to another forest.

A local user record is identified with the user **SID**. The command to migrate the user takes the original user account name and the new account name and indicates whether to validate the SID history.

In the case of a domain-to-domain transfer within the same forest, validation of SID history is recommended. When the MigrateUser command is issued, the Windows SharePoint Services front-end web server gets the new user SID from Active Directory from the new user account name, looks up the User Table for the record under the old user SID, and updates that row with the new user SID. In effect, this converts the user from one SID to another. When SID history is turned on, the new user token is examined to make sure that it contains the old user SID.

In the case of a forest-to-forest transfer, it is not possible to verify the SID history because the Active Directory forest is the boundary for security principals. In this case, the new SID is looked up from the new Active Directory forest (2), the User Table record that matches the old user SID is located, and the record is updated with the new SID.

2.9.2.5 Selecting Users and Groups from Active Directory

Active Directory is used by Windows SharePoint Services as the directory containing the list of users and groups that can be used for securing a container in Windows SharePoint Services. End users as well as administrators can look up users and groups, search for users and groups, and add one or more users and/or groups to the site.

Two basic functions are performed against Active Directory to select users and groups:

- Resolve a name or ID.
- Search for all matching records for a query.

In Windows SharePoint Services, a user is able to call the user and group selector in the security setting UI. The control contains a text box, a **Check Names** button, and a **Browse** button.

The typical scenario involves a user adding a user name or user ID into the text box and pressing the **Check Names** button. The front-end web server then sends the string to Active Directory and attempts to find a unique user or security group object that matches the text entered in the text box. If Active Directory locates a unique match, the object's SID is returned, and the Resolve call is considered successful. For an example of this operation, see the description of the Active Directory: People Picker Check Name UI in section 3.3.

If a unique object is not found, the Resolve call fails, and MAY return the list of matches for the query. The user interface marks the original text with a red underline, and when selected, will show all the possible matches returned from Active Directory in a drop-down list.

The other possibility is for the user to select the **Browse** button. This displays a pop-up window containing a search box and a results box. The user enters the search string and selects the **Go** button to issue the query. The search string is then sent to Active Directory, and all results are shown in the results box. The user is then able to select one or more of the results to add to the field. For an example of this operation, see the description of the Active Directory: People Picker Browse Display UI in section 3.2.

If no results are found, an informational message is displayed indicating that no results matched the query term.

The user and group selectors can be configured to select just users, just groups, or both.

2.9.2.6 Creating an Active Directory User Account

Windows SharePoint Services is often deployed on an intranet with Active Directory. In that case, Active Directory contains the list of users (for example, users of the corporation) who can potentially access the site. In this situation, it is impossible to invite someone who is not listed in Active Directory to view or participate in a Windows SharePoint Services site.

Windows SharePoint Services is also frequently used in the extranet to enable cross-company and/or cross-domain collaboration. In this case, it is not possible to pre-emptively place an exhaustive list of all users in all companies into Active Directory. In this situation, Windows SharePoint Services is deployed in a mode called Active Directory account creation mode. This allows Windows SharePoint Services to create an account for a user when that user first attempts to interact with a Windows SharePoint Services site.

For example, a Team Site administrator invites a partner using an email address, username@example.com. Windows SharePoint Services sends an email to that address inviting the user to access the resources at that Team Site. When the prospective user first arrives at the Windows SharePoint Services site, the user is asked to create a site account that identifies that user for future visits.

After the user completes the sign-up process, the Windows SharePoint Services front-end web server creates a user account in Active Directory, and stores the new user's name, email address, and other data. The location in Active Directory where these accounts are created, and the permissions to create accounts, are set up when Windows SharePoint Services is configured. For an example of this operation, see the description of the Active Directory: Account Creation New UI in section 3.1.

2.10 Additional Considerations

There are no additional considerations.

3 Examples

These examples describe in detail the process of communication between the various server components involved in the Windows SharePoint Services deployment. In conjunction with the technical protocol documents listed in section 2.2, these examples are intended to provide a comprehensive view of how Windows SharePoint Services front-end web servers communicate with end-user client, Active Directory domain controller (DC), and back-end database server systems.

3.1 Example 1: Active Directory: Account Creation New UI

This example describes the requests made when the user on an end-user client computer fills in the email address and **display name**, and then clicks the "OK" button on the Create Active Directory Account dialog to create a new Active Directory user account. The main member protocol used in this sequence is [MS-WSSFO] covering the stored procedures listed in the following steps. The sequence diagram has been broken into three figures because of size limitations. The three figures in this section represent a single sequence. This specific example is for Active Directory operations involving Windows SharePoint Services 3.0.

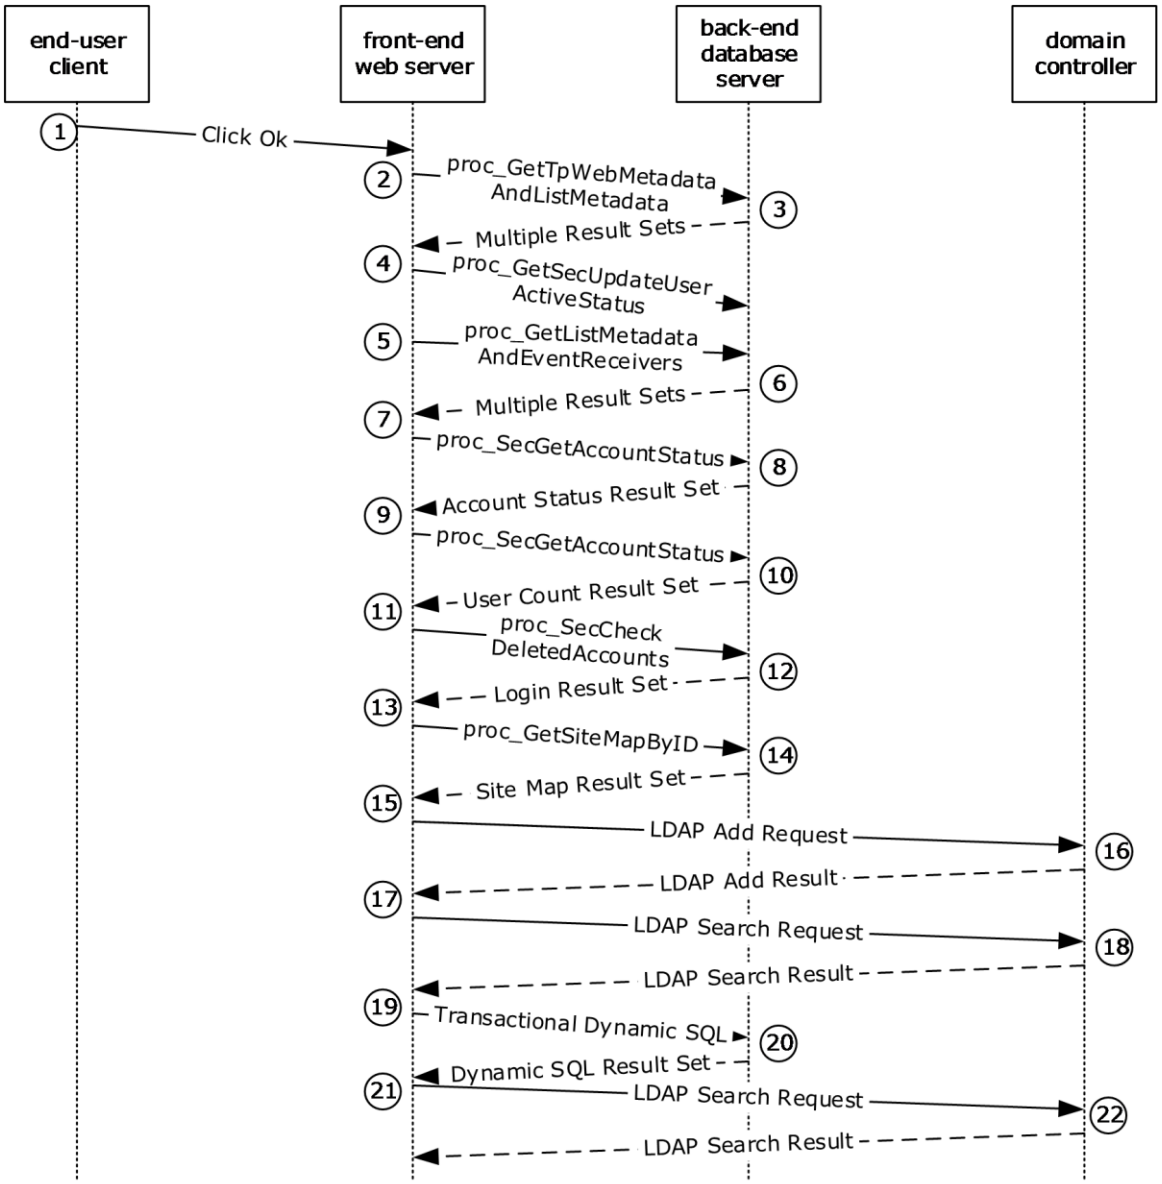

Figure 5: Account Creation New UI, steps 1 through 22

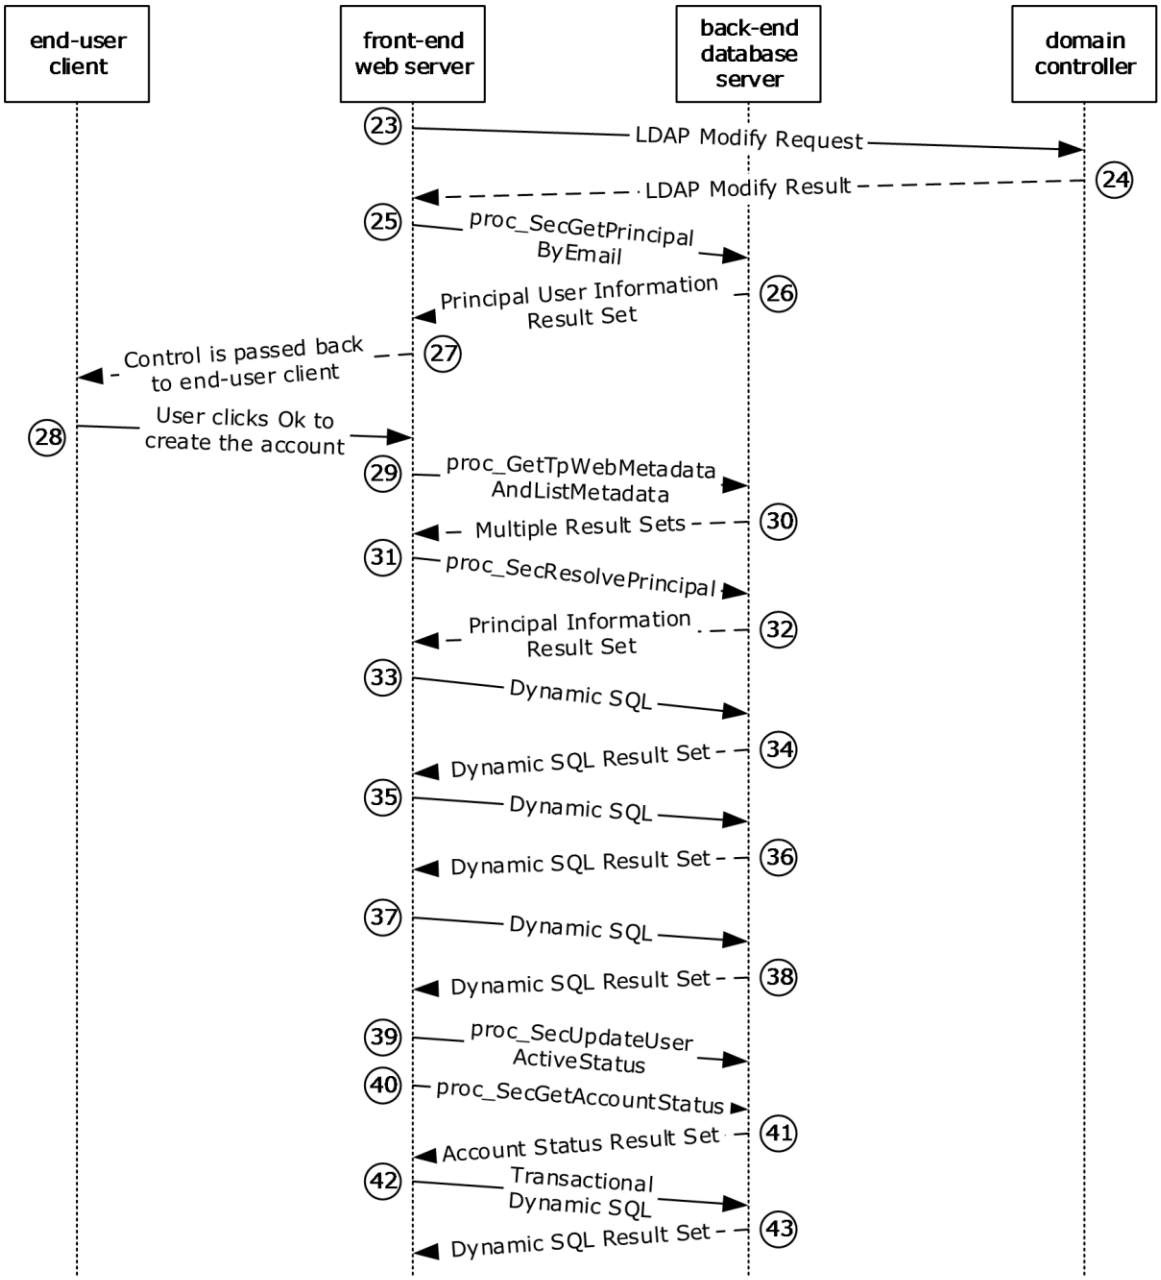

Figure 6: Account Creation New UI, steps 23 through 43

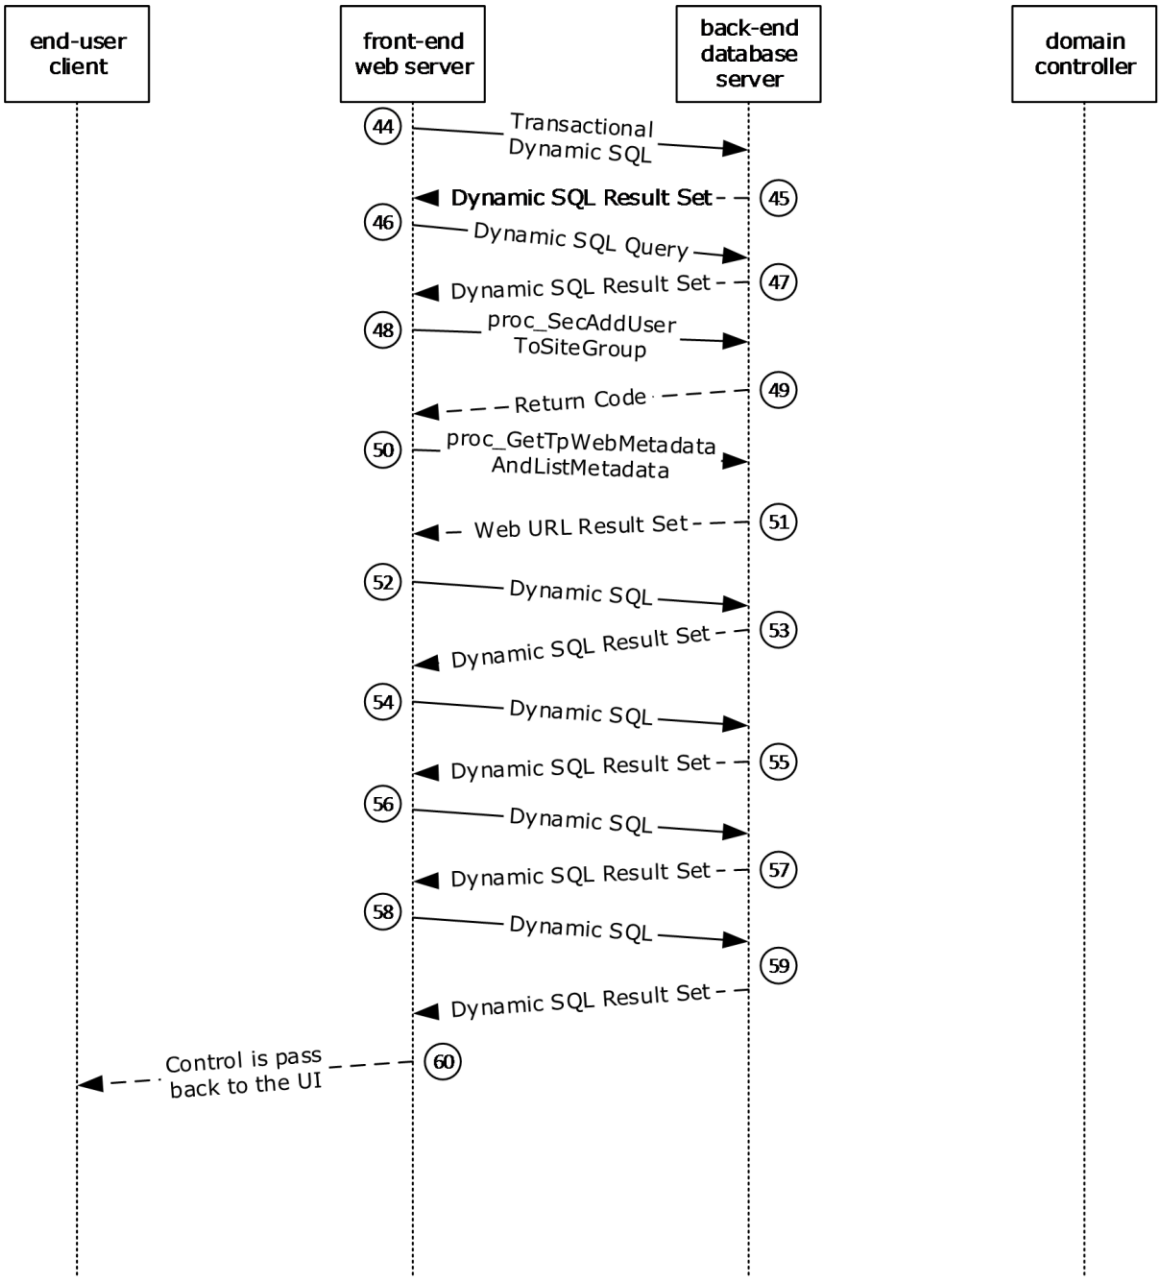

Figure 7: Account Creation New UI, steps 44 through completion

This scenario uses the Windows SharePoint Services user interface (UI) to create a new account for a user. It assumes that a user on an end-user client computer has selected the Add Users option on the people.aspx page under the **New** button. The user then clicks the **Create** button under the Users/Groups dialog, and fills in a valid email address and display name for a user who has never existed in this site collection. The following trace is then initiated when the user clicks the **OK** buttons on both the Create Active Directory Account page and the Add Users: Team Site page to add a new user to Active Directory.

The following actions happen:

1. The user clicks the **OK** button, which causes the end-user client to start the **UserAdd** process via a page to post back to the front-end web server to verify that the new user isn't currently defined.

2. The front-end web server first gathers information about the current environment by calling the proc_GetTpWebMetadataAndListMetadata stored procedure.

3. The back-end database server returns the following eight **result sets**:

   - **Web URL Result Set**, which returns the **store-relative URL** of the root of the site.
   - **Domain Group Cache Versions Result Set**, which returns information about the version numbers associated with the **domain group** map cache for this site.
   - **Domain Group Cache WFE Update Result Set**, which returns binary data needed to refresh the domain group map cache.
   - **Site Metadata Result Set**, which returns specialized site (2) metadata.
   - **Event Receivers Result Set**, which returns information about the **event receivers** defined for the site.
   - **Site Feature List Result Set**, which returns a list of default **feature identifiers** for the site collection that contains this site.
   - **Site Feature List Result Set**, which returns a list of feature identifiers for this site (2).
   - **Empty Result Set**, which is a **placeholder** set returned because the site has no cached navigation scope information.

4. The front-end web server calls the proc_SecUpdateUserActiveStatus stored procedure to update the list of active users for the site (2).

5. The front-end web server continues collecting information about the current user list by calling the proc_GetListMetadataAndEventReceivers stored procedure using the Tabular Data Stream (TDS) protocol, as specified in [MS-TDS].

6. The back-end database server returns the following two result sets:

   - **List Metadata Result Set**, which returns the metadata associated with this list.
   - **Event Receivers Result Set**, which returns information about the event receivers defined for this list (1).

7. The front-end web server checks for the existence of the new account by calling the proc_SecGetAccountStatus stored procedure.

8. The back-end database server returns an **Account Status Result Set** with zero rows, indicating that the email address has not yet been used in this site collection.

9. The front-end web server gathers further information about the current count of registered users by calling the proc_SecGetCurrentUsersCount stored procedure.

10. The back-end database server returns the **User Count Result Set** to indicate the number of users registered with this site collection and any quota information.

11. The front-end web server verifies that the UserID to be created does not exist in the deleted user list by calling the proc_SecCheckDeletedAccounts stored procedure.

12. The back-end database server returns the **Login Result Set** with zero results, indicating the UserID does not yet exist.

13. The front-end web server gathers further information about the current site (2) by calling the proc_GetSiteMapById stored procedure.

14. The back-end database server returns the **Site Map By Id Result Set** with one row of data on the current site (2).

15. The front-end web server makes an Lightweight Directory Access Protocol (LDAP) **AddRequest** to the domain controller (DC) to add a user object with the supplied information.

16. The DC sends an LDAP **AddResponse** indicating a successful insertion.

17. The front-end web server makes an LDAP **Search** request to the DC to confirm that the recently added user object exists.

18. The DC sends an LDAP **Search** response indicating that the user object exists.

19. The front-end web server builds a transactional dynamic SQL query that performs the following tasks:

    1. The query begins a new SQL **transaction**.

    2. The proc_SecAddUser stored procedure is executed to add the requesting user to Active Directory environment.

    3. In the event of any error, the transaction is rolled back, and the final select statement is returned.

    4. If the new user account does not already exist, it is added to the user list using the proc_AddListItem stored procedure.

    5. In the event of any error, the transaction is rolled back, and the final select statement is returned.

    6. The transaction is committed, and the final select statement is returned with the variables used to create the new User account.

20. The back-end database server returns a dynamic SQL result set, indicating that the new user account has been successfully added, and displaying the variables used when creating the new account.

21. The front-end web server then makes an LDAP **Search** request to the DC to request attributes for the recently added user object.

22. The DC sends an LDAP **Search** response indicating the additional user object attributes.

23. The front-end web server then makes an LDAP **Modify** request to the DC to modify the mail attribute for the added user object.

24. The DC sends an LDAP **Modify** response indicating the successful attribute change.

25. The front-end web server verifies the availability of the new account's email address by calling the proc_SecGetPrincipalByE-mail stored procedure.

26. The back-end database server returns the **Principal User Information Result Set**, containing information about the user associated with the specified email address.

27. Control is passed back to the end-user client on the Add Users page, with the new user listed in the Users/Groups dialog box.

28. The User clicks **OK** on the Add Users page to add the newly created user to the site group (section 2.9.1.5).

29. The front-end web server requests status data by calling the proc_GetTpWebMetadataAndListMetadata stored procedure.

30. The back-end database server returns the following 14 result sets:

    - **Web URL Result Set**, which returns the URL of the root of the site (2).
    - **Domain Group Cache Versions Result Set**, which returns information about the version numbers associated with the domain group map cache for this site (2).
    - **Domain Group Cache WFE Update Result Set**, which returns binary data needed to refresh the domain group map cache.
    - **Site Metadata Result Set**, which returns specialized site metadata.
    - **Event Receivers Result Set**, which returns information about the event receivers defined for this site (2).
    - **Site Category Result Set**, which returns categories of this site (2).
    - **Site Metainfo Result Set**, which returns the specialized site metadata.
    - **Site Feature List Result Set**, which returns a list of default feature identifiers for the site collection that contains this site (2).
    - **Site Feature List Result Set**, which returns a list of feature identifiers of this site (2).
    - **Empty Result Set**, which is a placeholder set returned because the site (2) has no cached navigation scope information.
    - **List Metadata Result Set**, which returns the metadata associated with the specified document list.
    - **NULL Unique Permissions Result Set**, which is a placeholder set returned because the list has no individual list permissions.
    - **Event Receivers Result Set**, which returns information about the event receivers defined for this document list (1).
    - **List Web Parts Result Set**, which returns information about the list **Web Parts** defined for this document list.

31. The front-end web server requests information about the new account by calling the proc_SecResolvePrincipal stored procedure.

32. The back-end database server returns the **Principal Information Result Set** with a single row containing basic information about the user.

33. The front-end web server creates a dynamic SQL query, which selects information from the UserData view joined with the Docs view.

34. The back-end database server returns a dynamic SQL result set, which contains one row of data with the new user account information.

35. The front-end web server creates a dynamic SQL query, which selects information from Sec_SiteGroupsView.

36. The back-end database server returns a dynamic SQL result set with all site group membership levels.

37. The front-end web server creates a dynamic SQL query to check the requesting user permissions for this activity by calling the proc_SecGetUserPermissionOnGroup stored procedure.

38. The back-end database server returns a dynamic SQL result set representing the permission levels of the calling user.

39. The front-end web server calls the proc_SecUpdateUserActiveStatus stored procedure to update the list of active users for the site (2).

40. The front-end web server requests the account status for the new account by calling the proc_SecGetAccountStatus stored procedure.

41. The back-end database server returns the **Account Status Result Set**.

42. The front-end web server builds a transactional dynamic SQL query which performs the following tasks:

    1. The query begins a new SQL transaction (3).

    2. The proc_SecAddUser stored procedure is executed to add the requesting user to the appropriate security groups in the Active Directory environment.

    3. In the event of any error, the transaction (3) is rolled back, and the final select statement is returned.

    4. If the new user account does not already exist, it is added to the user list using the proc_AddListItem stored procedure.

    5. In the event of any error, the transaction is rolled back and the final select statement is returned.

    6. The transaction is committed, and the final select statement is returned with the variables used to create the new user account.

43. The back-end database server returns a dynamic SQL result set indicating the success of the add procedures and the variables used in the new account.

44. The front-end web server builds a transactional dynamic SQL query, which performs the following tasks:

    1. The query begins a new SQL transaction (3).

    2. The proc_AddListItem stored procedure is executed to add the new user account to the appropriate Windows SharePoint Services list.

    3. In the event of any error, the transaction (3) is rolled back and the final select statement is returned.

    4. The new user account is updated, using the proc_UpdateListItem stored procedure with additional user data as necessary.

    5. In the event of any error, the transaction is rolled back and the final select statement is returned.

    6. The transaction is committed, and the final select statement is returned with the variables used to create the new user account.

45. The back-end database server returns a dynamic SQL result set with information about the newly created user.

46. The front-end web server creates a dynamic SQL query to either add the user data by calling the proc_AddListItem stored procedure, or to update the user data by calling the proc_UpdateListItem stored procedure.

47. The back-end database server returns the following two result sets:

    - **Item Update Result Set**, which returns pertinent information about the update.

- **Dynamic SQL Result Set**, which returns the output status value from the update.

48. The front-end web server can now add the new user account to the appropriate site group by calling the proc_SecAddUserToSiteGroup stored procedure.

49. The back-end database server responds with a **return code**, but no result sets are returned.

50. The front-end web server requests further status data by calling the proc_GetTpWebMetadataAndListMetadata stored procedure.

51. The back-end database server returns the following 14 result sets:

    - **Web URL Result Set**, which returns the URL of the root of the site (2).
    - **Domain Group Cache Versions Result Set**, which returns information about the version numbers associated with the domain group map cache for this site (2).
    - **Domain Group Cache WFE Update Result Set**, which returns binary data needed to refresh the domain group map cache.
    - **Site Metadata Result Set**, which returns specialized site metadata.
    - **Event Receivers Result Set**, which returns information about the event receivers defined for this site (2).
    - **Site Category Result Set**, which returns categories of this site (2).
    - **Site Metainfo Result Set**, which returns the specialized site metadata.
    - **Site Feature List Result Set**, which returns a list of default feature identifiers for the site collection that contains this site (2).
    - **Site Feature List Result Set**, which returns a list of feature identifiers of this site (2).
    - **Empty Result Set**, which is a placeholder set.
    - **List Metadata Result Set**, which returns the metadata associated with the specified document list (1).
    - **NULL Unique Permissions Result Set**, which is a placeholder set.
    - **Event Receivers Result Set**, which returns information about the event receivers defined for the document list (1).
    - **List Web Parts Result Set**, which returns information about the list (1) Web Parts defined for this document list (1).

52. The front-end web server creates a dynamic SQL query selecting all information from the Sec_SiteGroupsView view for this site (2) to generate the final website (2) display.

53. The back-end database server returns a dynamic SQL result set with all site group membership levels.

54. The front-end web server creates a dynamic SQL query, selecting information from the UserData view, Docs view, and AllUserData view for details on the website (2) display.

55. The back-end database server returns a dynamic SQL result set of user data as it applies to the current display.

56. The front-end web server creates a dynamic SQL request for user permission information by calling the proc_SecGetUserPermissionOnGroup stored procedure.

57. The back-end database server returns a dynamic SQL result set representing the permission levels of the calling user.

58. The front-end web server creates a dynamic SQL request for information from the UserData view and Docs view for website (2) display.

59. The back-end database server returns a dynamic SQL result set of user data as it applies to the current display.

60. Control is then returned to the UI.

3.2 Example 2: Active Directory: People Picker Browse Display UI

This example describes the requests that are made when a search for a valid Active Directory user is made from the end-user client computer by entering a search string that matches a user's display name, and when that user is located, that user is added to the current site. The main member protocol used in this sequence is [MS-WSSFO] covering the stored procedures listed in the steps. The sequence diagram has been broken into three figures because of size limitations. The three figures in this section represent a single sequence. This specific example is for Active Directory operations involving Windows SharePoint Services 3.0.

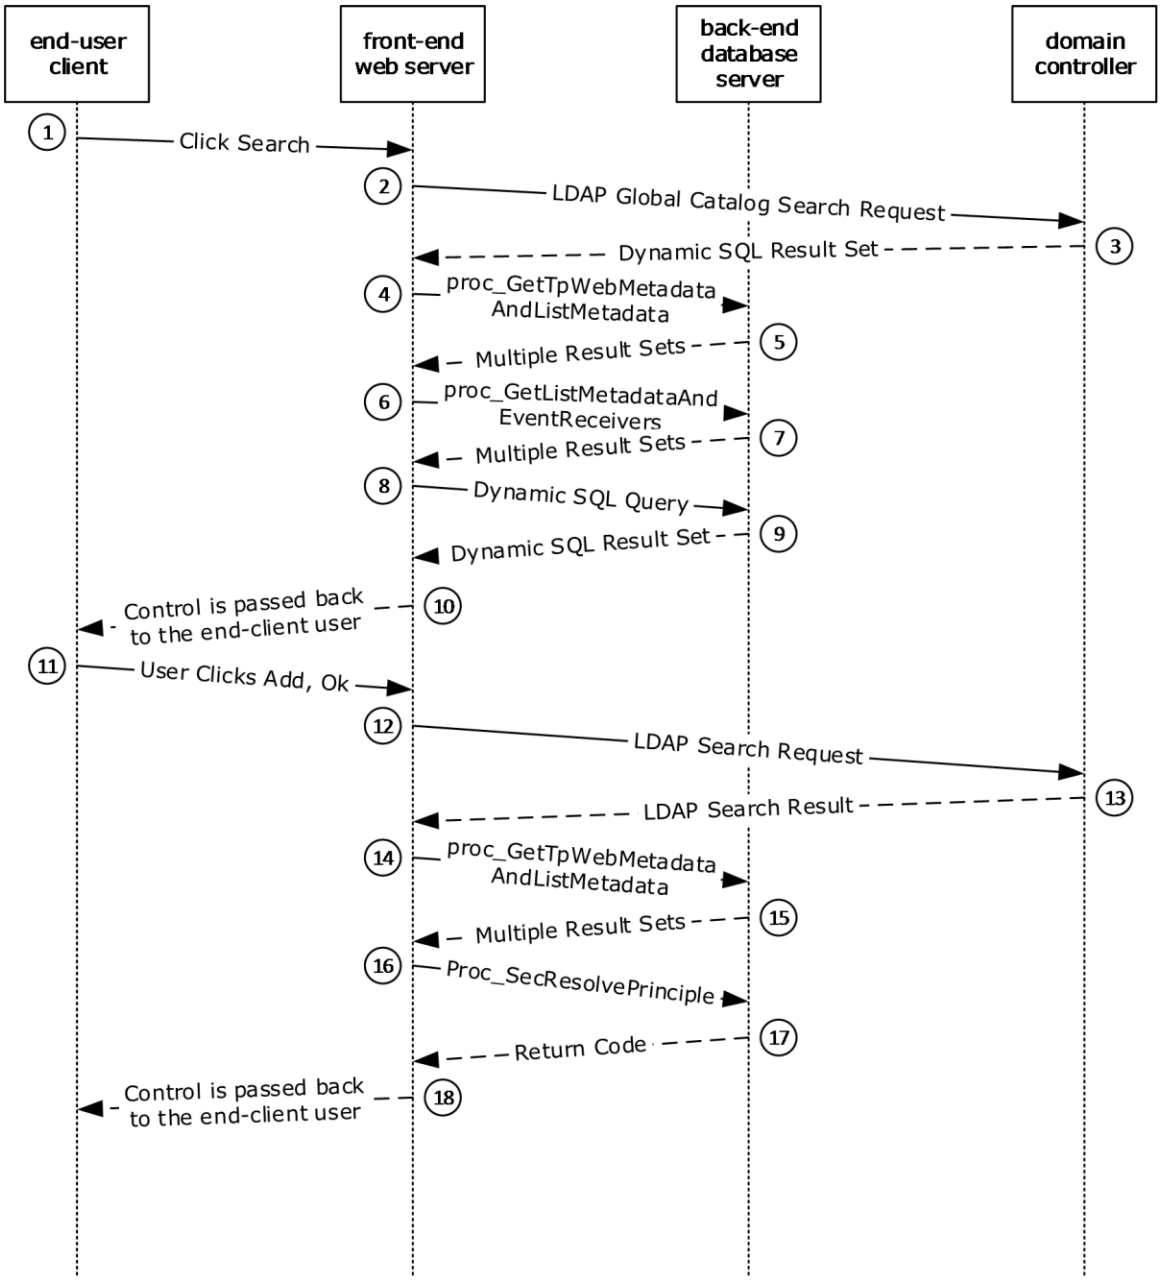

Figure 8: People Picker Browse Display UI, steps 1 through 18

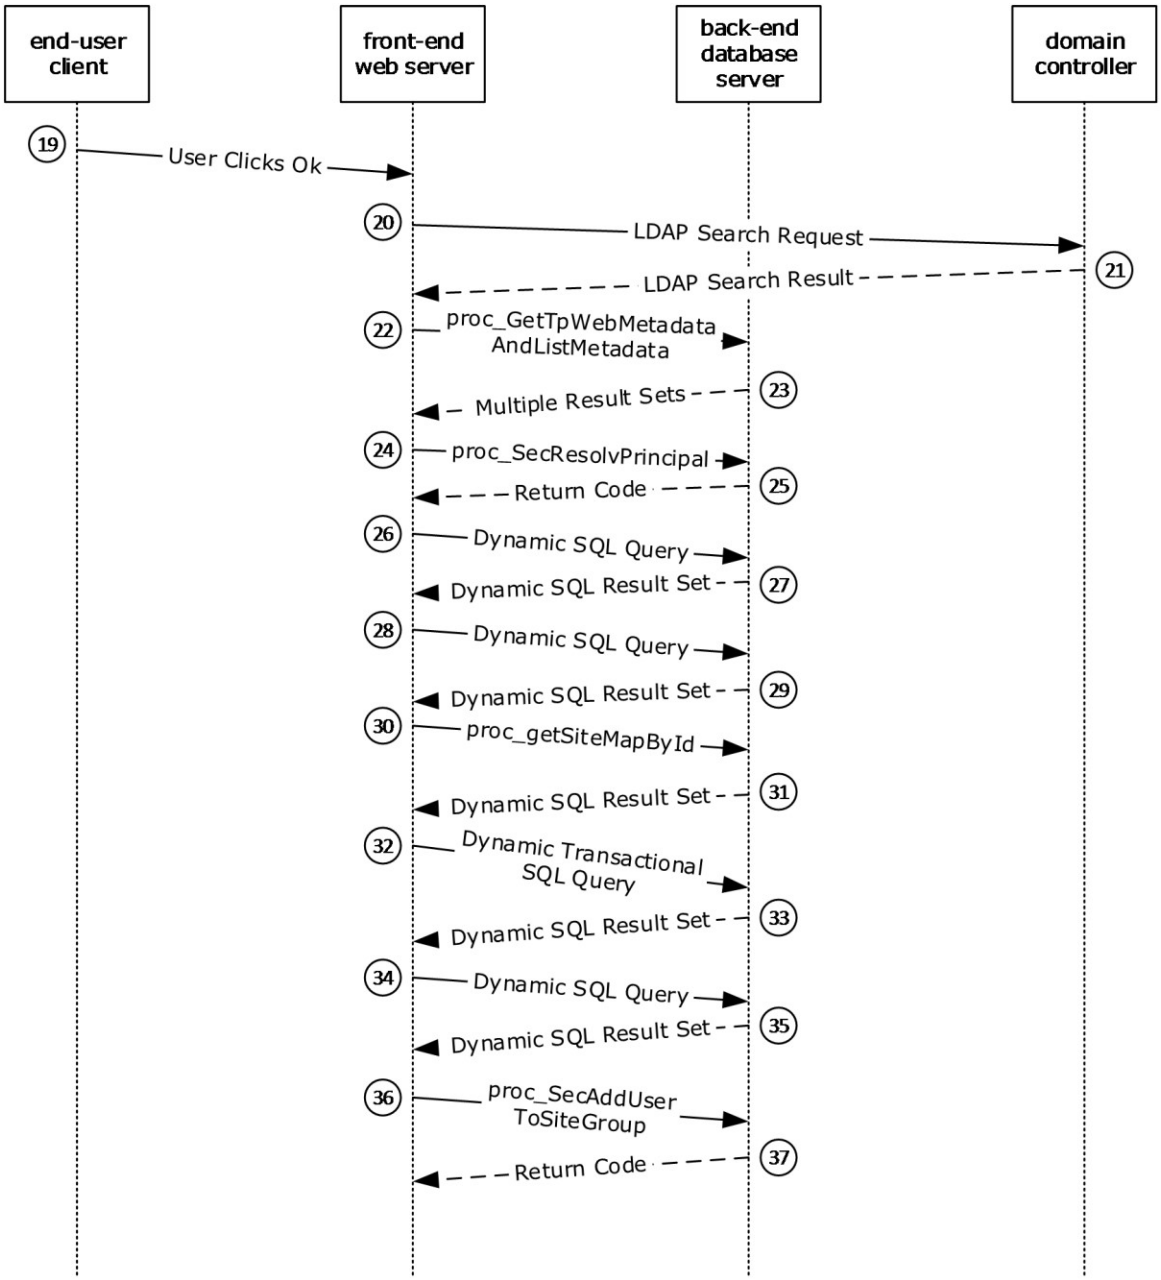

Figure 9: People Picker Browse Display UI, steps 19 through 37

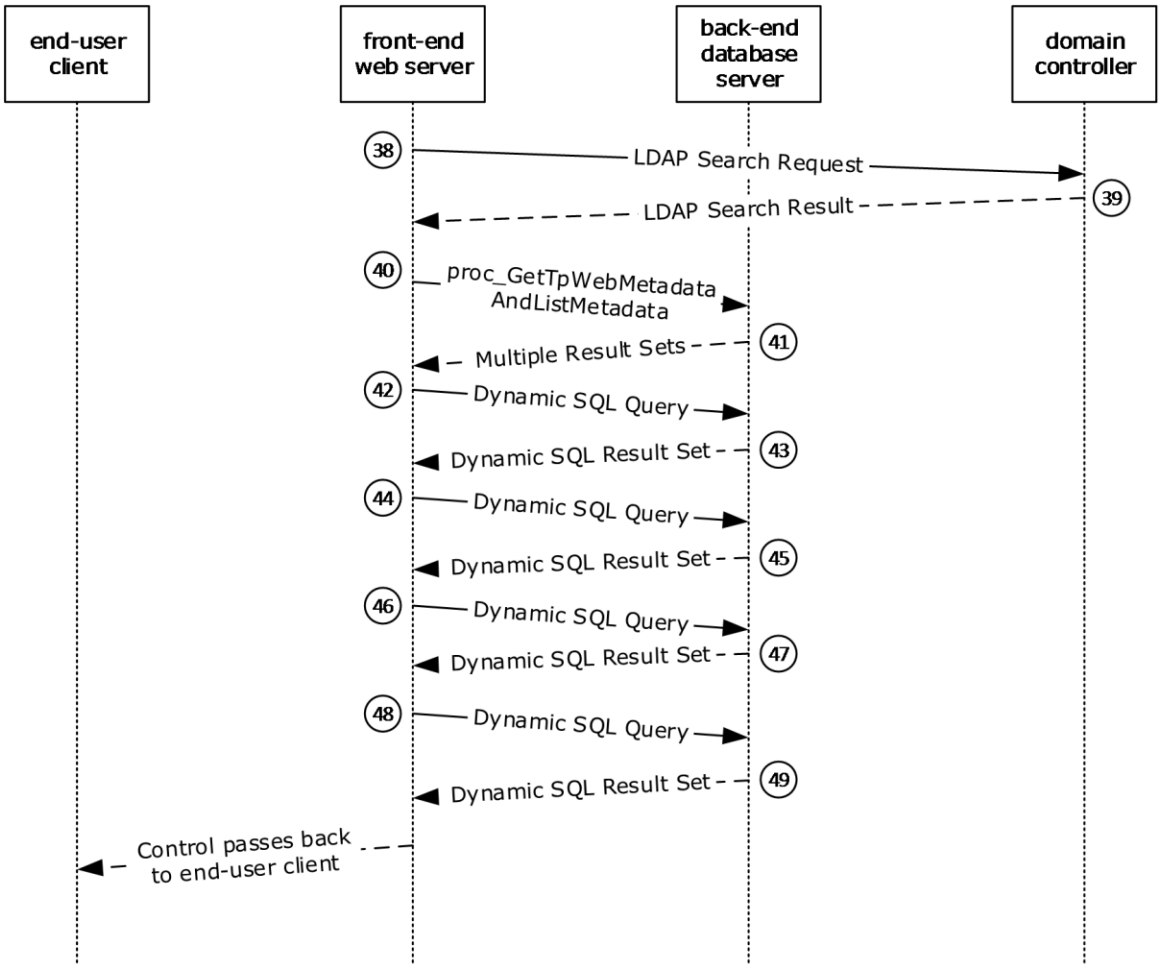

Figure 10: People Picker Browse Display UI, steps 38 through completion

This scenario is initiated from the "Select People and Groups" dialog. A user enters a search string in the "Find" text field and then clicks the search icon. For the sake of simplicity, it is assumed that the user has Add privileges for the current site group (section 2.9.1.5).

The following actions happen:

1. The end-user client first sends a request to the front-end web server to search for the desired User display name.

2. The front-end web server sends a Lightweight Directory Access Protocol (LDAP) **Global Catalog Search** request to the domain controller (DC) asking for any match in the whole subtree for user objects or **group objects** with attributes that contain the search string (a wildcard search version of the user display name) in one of the following attributes:

   - **User objects**: "name", "displayName", "cn", "sn", "SamAccountName", "mail", **Simple Mail Transfer Protocol (SMTP)** or Session Initiation Protocol (SIP) "proxyAddresses" attributes.

   - **Group objects**: "name", "displayName", "cn", or "SamAccountName" attributes.

3. The DC responds with an LDAP **Global Catalog Search** response containing both user objects and group objects that match the search string.

4. The front-end web server initializes information about the site and its users by calling the proc_GetTpWebMetadataAndListMetadata stored procedure using the Tabular Data Stream (TDS) protocol, as specified in [MS-TDS].

5. The back-end database server returns five result sets:

   - **Web URL Result Set**, which returns the store-relative URL of the root of the site.
   - **Domain Group Cache Versions Result Set**, which returns information about the version numbers associated with the domain group map cache for this site.
   - **Domain Group Cache WFE Update Result Set**, which returns information to be used in recomputing the domain group map cache for the site (2).
   - **Site Metadata Result Set**, which returns specialized site metadata.
   - **Event Receivers Result Set**, which returns information about the event receivers defined for the site.

6. The front-end web server continues collecting information about the site's user list by calling the proc_GetListMetadataAndEventReceivers stored procedure.

7. The back-end database server returns the following four result sets:

   - **List Metadata Result Set**, which returns the permissions associated with the user list.
   - **NULL Unique Permissions Result Set**, which indicates that unique permissions do not exist for the list.
   - **List Event Receivers Result Set**, which is empty because there are no event receivers defined for this list (1).
   - **List Web Parts Result Set**, which contains information about the **list view** pages.

8. The front-end web server creates a dynamic SQL query that searches for the submitted search string in the user information list, looking for a match in the display name, account name or email address columns.

9. The back-end database server returns one empty dynamic SQL result set, indicating that a match was not found.

10. The front-end web server displays the display name received from the DC as a candidate for selection.

11. The end user clicks **Add**, then **OK**. The end-user client closes the dialog and redirects the user to the User Information List web page.

12. The front-end web server negotiates authentication with the DC and then sends an LDAP **Search** request to the DC for an object that has a security identifier (SID) attribute equal to the value obtained from the DC in Step 3.

13. The DC sends an LDAP **Search** result containing the attributes of the Active Directory user object.

14. The front-end web server again initializes by gathering information about the site (2) by calling the proc_GetTpWebMetadataAndListMetadata stored procedure.

15. The back-end database server returns five result sets:

   - **Web URL Result Set**, which contains the store-relative URL of the root of the site (2).
   - **Domain Group Cache Versions Result Set**, which contains information about the version numbers associated with the domain group map cache for this site (2).

- **Domain Group Cache WFE Update Result Set**, which contains information to be used in recomputing the domain group map cache for the site (2).
- **Site Metadata Result Set**, which contains site metadata.
- **Event Receivers Result Set**, which contains information about the event receivers defined for the site (2).

16. The front-end web server sends a request to the back-end database server to find security principals that might have **login name**, display name, or email address information matching the user account name returned from the DC. It does so by calling the proc_SecResolvePrincipal stored procedure.
17. The back-end database server responds with a return code, but no result sets are returned, indicating that no matches were found.
18. The front-end web server renders the name as resolved.
19. The end user clicks **OK** on the Add Users page, sending a request to the front-end web server to add the user to the site and site group.
20. The front-end web server negotiates authentication with the DC, and then sends an LDAP **Search** request to the DC for an object that has a SID attribute equal to the value obtained from the DC in Step 3.
21. The DC sends an LDAP **Search** result containing the attributes of the Active Directory user object.
22. The front-end web server initializes again by calling the proc_GetTpWebMetadataAndListMetadata stored procedure.
23. The back-end database server returns the following 14 result sets:
    - **Web URL Result Set**, which contains the URL of the site (2).
    - **Domain Group Cache Versions Result Set**, which contains information about the version numbers associated with the domain group map cache for this site (2).
    - **Domain Group Cache WFE Update Result Set**, which contains binary data needed to refresh the domain group map cache.
    - **Site Metadata Result Set**, which contains site metadata.
    - **Event Receivers Result Set**, which contains information about the event receivers that are defined for this site (2).
    - **Site Category Result Set**, which contains the categories of this site (2).
    - **Site Metainfo Result Set**, which contains the specialized site metadata.
    - **Site Feature List Result Set**, which contains the list of default feature identifiers for the site collection that contains this site.
    - **Site Feature List Result Set**, which contains the list of feature identifiers of this site (2).
    - **Empty Result Set**, which is a placeholder set.
    - **List Metadata Result Set**, which contains the metadata associated with the specified document list.
    - **NULL Unique Permissions Result Set**, which indicates that there are no special permissions set on the user information list.

- **Event Receivers Result Set**, which contains information about the event receivers defined for the document list (1).
- **List Web Parts Result Set**, which contains information about the list view pages defined for the user information list (1).

24. The front-end web server sends a request to resolve the selected user names by calling the proc_SecResolvePrincipal stored procedure.
25. The back-end database server responds with a return code, but no result sets are returned, indicating that the user was not found.
26. The front-end web server creates a dynamic SQL query that selects information from the Sec_SiteGroupsView.
27. The back-end database server returns a dynamic SQL result set with all site group membership levels signifying the owner of all groups.
28. The front-end web server builds a dynamic SQL query to determine whether the current user has permission to add a user to the group. It does this by calling the proc_SecGetUsersPermissionsOnGroup stored procedure.
29. The back-end database server returns one dynamic SQL result set, which contains one record for the current group, indicating that the current user does not directly have permission to add a user to the group, and is not the owner of the group.
30. The front-end web server requests the site map by calling the proc_getSiteMapById stored procedure.
31. The back-end database server returns the **Site Map by Id Result Set**.
32. The front-end web server builds a dynamic transactional SQL Query to add the user to the site collection. The following actions happen:
    1. The transaction begins.
    2. An attempt to add a user to the UserInfo table is performed by calling the proc_SecAddUser stored procedure.
    3. If adding the user succeeded, then an attempt to add a person list item to the User Information List is performed. It does so by calling the proc_AddListItem stored procedure.
    4. If either adding the user to the site collection or adding the list item to the User Information List failed, then the transaction is rolled back; otherwise, the transaction is committed.
33. One result is returned from the back-end database server, containing the return code and information about the added user.
34. The front-end web server constructs a dynamic SQL query, selecting full user information about the added user.
35. The back-end database server returns a dynamic SQL result set with the requested information.
36. The front-end web server sends a request to the back-end database server to add the user to the current site group by calling the proc_SecAddUserToSiteGroup stored procedure.
37. The back-end database server responds with a return code, but no result sets are returned.
38. The front-end web server negotiates authentication with the DC, and then sends an LDAP **Search** request to the DC for an object that has a SID attribute equal to the value obtained from the DC in Step 3.

39. The DC sends an LDAP **Search** result containing the attributes of the Active Directory user object.

40. The front-end web server again initializes its information about the site (2) by calling the proc_GetTpWebMetadataAndListMetadata stored procedure.

41. The back-end database server returns the following 14 result sets:

    - **Web URL Result Set**, which returns the URL of the root of the site (2).
    - **Domain Group Cache Versions Result Set**, which returns information about the version numbers associated with the domain group map cache for this site (2).
    - **Domain Group Cache WFE Update Result Set**, which returns binary data needed to refresh the domain group map cache.
    - **Site Metadata Result Set**, which returns specialized site metadata.
    - **Event Receivers Result Set**, which returns information about the event receivers defined for this site (2).
    - **Site Category Result Set**, which returns the categories of the site (2).
    - **Site Metainfo Result Set**, which returns the specialized site metadata.
    - **Site Feature List Result Set**, which returns the list of default feature identifiers for the site collection that contains this site (2).
    - **Site Feature List Result Set**, which returns the list of feature identifiers of this site (2).
    - **Empty Result Set**, which is a placeholder set.
    - **List Metadata Result Set**, which returns the metadata associated with the specified document list.
    - **NULL Unique Permissions Result Set**, which is a placeholder set.
    - **Event Receivers Result Set**, which returns information about the event receivers defined for the document list (1).
    - **List Web Parts Result Set**, which returns information about the list (1) Web Parts defined for this document list (1).

42. The front-end web server creates a dynamic SQL query that selects information from the Sec_SiteGroupsView view.

43. The back-end database server returns a dynamic SQL result set with all site group membership levels, signifying the owner of all groups.

44. The front-end web server builds a dynamic SQL query to obtain updated information about the site group to which the user was added.

45. The back-end database server returns one dynamic SQL result set containing information about the site group.

46. The front-end web server builds a dynamic Query to determine whether the current user has permission to add a user to the group. It does this by calling the proc_SecGetUsersPermissionsOnGroup stored procedure.

47. The back-end database server returns one dynamic SQL result set, which contains one record for the current group, indicating that the current user does not directly have permission to add a user to the group and is also not the owner of the group.

48. The front-end web server builds a dynamic SQL query to obtain more user information for the site group to which the user has been added.

49. The back-end database server returns one dynamic SQL result set of information about the newly added user.

50. Control is passed back to the end-user client.

3.3 Example 3: Active Directory: People Picker Check Name UI

This example describes the requests made when the user is adding a new member to a SharePoint list and uses the Check Names function to confirm the existence of the new member in Active Directory. The main member protocol used in this sequence is [MS-WSSFO] covering the stored procedures listed in the steps. The sequence diagram has been broken into three figures because of size limitations. The three figures in this section represent a single sequence. This specific example is for Active Directory operations involving Windows SharePoint Services 3.0.

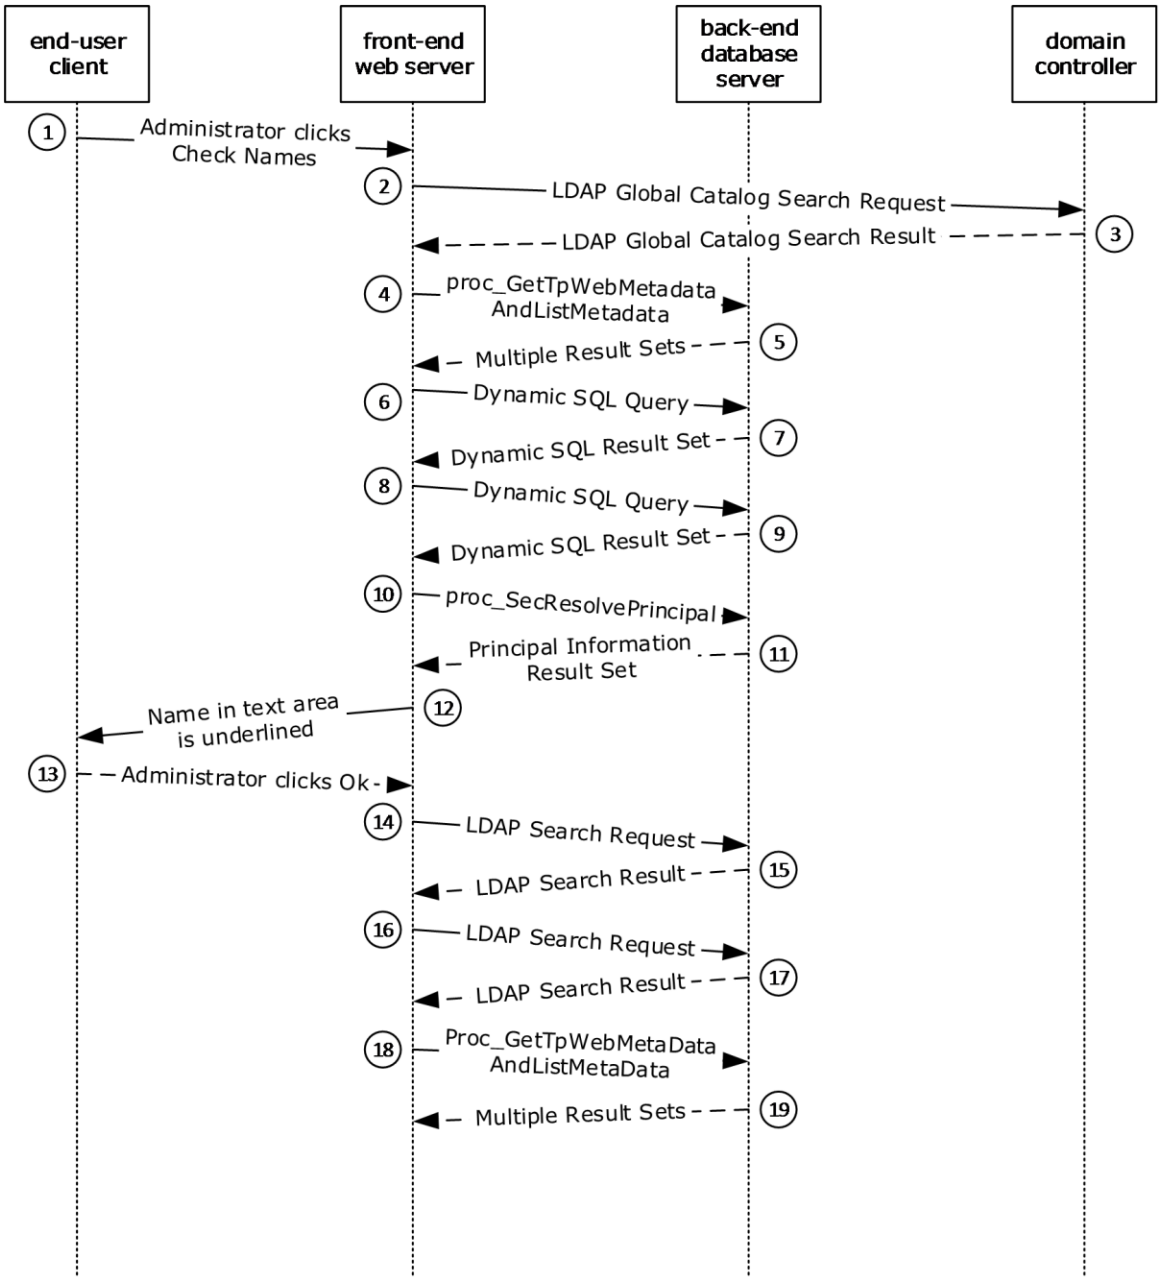

Figure 11: People Picker Check Name UI, steps 1 through 19

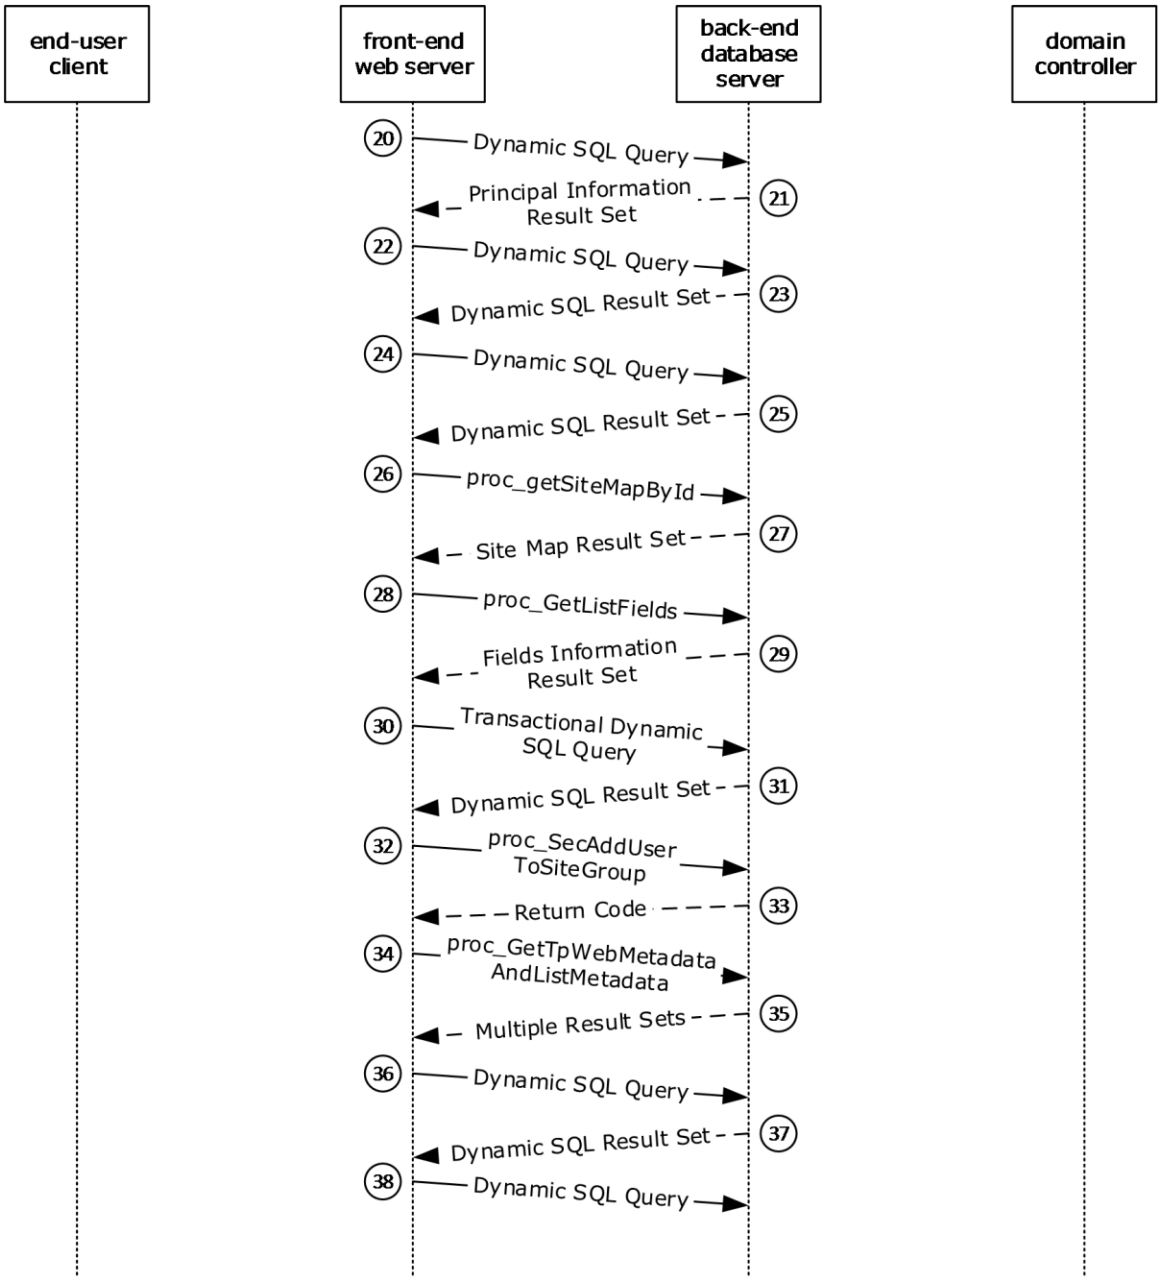

Figure 12: People Picker Check Name UI, steps 20 through 38

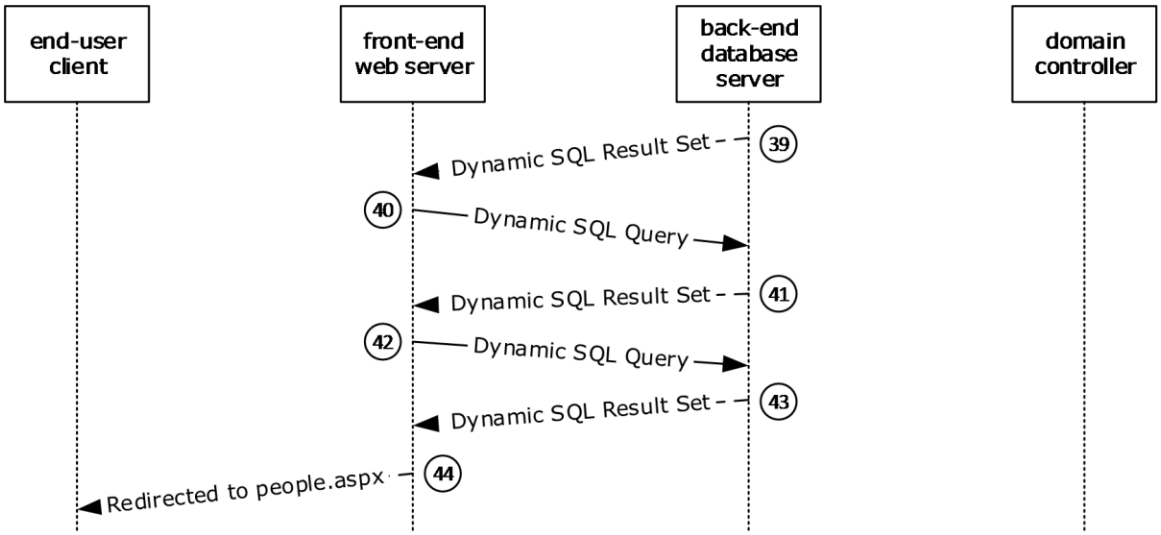

Figure 13: People Picker Check Name UI, steps 39 through completion

This scenario is initiated when the user clicks the Check Names icon. The following assumptions are made for the purposes of this example:

- The Windows SharePoint Services site is already established.
- The user making the request has permission to do so.
- The name to be added already exists in Active Directory.

The following actions happen:

1. From the end-user client, the user clicks the "Check Names" icon.
2. The front-end web server sends an Lightweight Directory Access Protocol (LDAP) **Global Catalog Search** request to the domain controller (DC), asking for any match in the whole subtree for user or group objects with attributes that contain the search string (a wildcard search version of the user display name) in one of the following attributes:
    - **User objects**: "name", "displayName", "cn", "samAccountName", "mail", Simple Mail Transfer Protocol (SMTP) or Session Initiation Protocol (SIP) "proxyAddresses" attributes.
    - **Group objects**: "name", "displayName", "cn", or "samAccountName" attributes.
3. The DC responds with an LDAP **Global Catalog Search** response containing both user objects and group objects that match the search string.
4. The front-end web server first collects metadata for the site and the document list by calling the proc_GetTpWebMetaDataAndListMetaData stored procedure using the Tabular Data Stream (TDS) protocol, as specified in [MS-TDS].
5. The back-end database server returns 14 result sets:
    - **Web Url Result Set**, which returns the store-relative URL of the root of the site.
    - **Domain Group Cache Versions Result Set**, which returns information about the version numbers associated with the domain group map cache for the site.
    - **Domain Group Cache WFE Update Result Set**, which returns the binary data needed to refresh the domain group map cache.

- **Site Metadata Result Set**, which returns specialized site (2) metadata.
- **Event Receivers Result Set**, which returns information about the event receivers defined for the site.
- **Site Category Result Set**, which returns the categories of the site (2).
- **Site MetaInfo Result Set**, which returns the specialized site (2) metadata.
- **Site Feature List Result Set**, which returns a list of feature identifiers for the site collection that contains this site.
- **Site Feature List Result Set**, which returns a list of feature identifiers for the site (2).
- **Empty Result Set**, which is returned because the **METADATA_WEB** and **METADATA_WEB_NAVSTRUCT** flags are set and the site has no cached scope information.
- **List Metadata Result Set**, which returns metadata associated with the document list (1).
- **NULL Unique Permissions Result Set**, which is returned because the **METADATA_PREFETCH_SCOPES** flag is set, the **METADATA_LISTMETADATA_NOFETCH** flag has not been set, and permissions for the document list do not exist.
- **Event Receivers Result Set**, which returns information about the event receivers defined for the document list (1).
- **List Web Parts Result Set**, which returns information about the list Web Parts defined for the document list.

6. The front-end web server builds a dynamic SQL query to retrieve information from the Sec_SiteGroupsView view. This query is sent to the SQL server using TDS.
7. The back-end database server returns a dynamic SQL result set containing information about the Team Site owners, the Team Site visitors, and the Team Site members (3).
8. The front-end web server builds a dynamic SQL query regarding which permissions the requesting user has within the site group (section 2.9.1.5). It does this by calling the proc_SecGetUserPermissionOnGroup stored procedure using TDS.
9. The back-end database server returns a dynamic SQL result set consisting of the output variables of the stored procedure.
10. The front-end web server then retrieves information about the user to be added by calling the proc_SecResolvePrincipal stored procedure using TDS.
11. The back-end database server returns the **Principal Information Result Set**, consisting of a single row of information about the user to be added.
12. When the information about the user to be added has been confirmed, the user's name in the text area is underlined.
13. The user clicks the "OK" button.
14. The front-end web server negotiates authentication with the DC and then sends an LDAP **Search** request to the DC for an object that has a security identifier (SID) attribute equal to the value obtained from the DC earlier.
15. The DC sends an LDAP **Search** result containing the attributes of the Active Directory user object.

16. The front-end web server again negotiates authentication with the DC, and then sends an LDAP **Search** request to the DC for an object that has a SID attribute equal to the value obtained from the DC in Step 3.

17. The DC sends an LDAP **Search** result containing the attributes of the Active Directory user object.

18. The front-end web server then collects metadata for the site (2) and the document list (1) by calling the proc_GetTpWebMetaDataAndListMetaData stored procedure using TDS.

19. The back-end database server returns 14 result sets:

    - **Web Url Result Set**, which returns the store-relative URL of the root of the site (2).
    - **Domain Group Cache Versions Result Set**, which returns information about the version numbers associated with the domain group map cache for the site (2).
    - **Domain Group Cache WFE Update Result Set**, which returns the binary data needed to refresh the domain group map cache.
    - **Site Metadata Result Set**, which returns specialized site (2) metadata.
    - **Event Receivers Result Set**, which returns information about the event receivers defined for the site (2).
    - **Site Category Result Set**, which returns the categories of the site (2).
    - **Site MetaInfo Result Set**, which returns the specialized site (2) metadata.
    - **Site Feature List Result Set**, which returns a list of feature identifiers for the site collection that contains this site (2).
    - **Site Feature List Result Set**, which returns a list of feature identifiers for the site (2).
    - **Empty Result Set**, which is returned because the **METADATA_WEB** and **METADATA_WEB_NAVSTRUCT** flags are set, and the site has no cached scope information.
    - **List Metadata Result Set**, which returns metadata associated with the document list (1).
    - **NULL Unique Permissions Result Set**, which is returned because the **METADATA_PREFETCH_SCOPES** flag is set, the **METADATA_LISTMETADATA_NOFETCH** flag has not been set, and permissions for the document list do not exist.
    - **Event Receivers Result Set**, which returns information about the event receivers defined for the document list (1).
    - **List Web Parts Result Set**, which returns information about the list (1) Web Parts defined for the document list (1).

20. The front-end web server then retrieves information about the user to be added by calling the proc_SecResolvePrincipal stored procedure using TDS.

21. The back-end database server returns the **Principal Information Result Set**, consisting of a single row of information about the user to be added.

22. The front-end web server builds a dynamic SQL query to retrieve information from the Sec_SiteGroupsView view. This query is sent to the SQL server using TDS.

23. The back-end database server returns a dynamic SQL result set containing information about the Team Site owners, the Team Site visitors, and the Team Site members (3).

24. The front-end web server builds a dynamic SQL query regarding what permissions the requesting user has within the site group. It does this by calling the proc_SecGetUserPermissionOnGroup stored procedure using TDS.

25. The back-end database server returns a dynamic SQL result set consisting of the output variables of the proc_SecGetUserPermissionOnGroup stored procedure.

26. The front-end web server then retrieves the database and URL mapping information for the site collection. It does this by calling the proc_getSiteMapById stored procedure from the Windows SharePoint Services configuration database using TDS.

27. The back-end database server returns the **Site Map by Id Result Set**.

28. The front-end web server then retrieves the mapping of **fields** in the document list by calling the proc_GetListFields stored procedure using TDS.

29. The back-end database server returns the **Fields Information Result Set**, consisting of a single row containing a single column of a Windows SharePoint Services implementation-specific version string followed by an **XML** representation of the field definitions.

30. The front-end web server builds a transactional dynamic SQL query to add an entry for the new user to the list of user information stored in the back-end database server. This query is sent to the SQL server using TDS. On the SQL server, the following actions occur:

    1. The query begins a new SQL transaction.

    2. The query attempts to add the new user to the list of user information in the back-end database server by calling the proc_SecAddUser stored procedure using TDS.

    3. The query rolls back the SQL transaction if the proc_SecAddUser stored procedure's return code is not "0", or it checks to see if the new user's UserId exists in the UserData table.

    4. If the user's UserId is not found in the UserData table, the query attempts to add the list item to the document list by calling the proc_AddListItem stored procedure using TDS.

    5. The query rolls back the SQL transaction if the proc_AddListItem stored procedure's Return Code is not "0", or it commits the transaction if the proc_SecAddUser's return code is "0" and proc_AddListItem's return code (if it runs) is "0".

31. The back-end database server returns a dynamic SQL result set that consists of information about the new user.

32. The front-end web server then attempts to add the new user to the site group. It does this by calling the proc_SecAddUserToSiteGroup stored procedure using TDS.

33. The back-end database server responds with a return code, but no result sets are returned.

34. The front-end web server then collects metadata for the site (2) and the document list (1) by calling the proc_GetTpWebMetaDataAndListMetaData stored procedure using TDS.

35. The back-end database server returns 14 result sets:

    - **Web Url Result Set**, which returns the store-relative URL of the root of the site (2).

    - **Domain Group Cache Versions Result Set**, which returns information about the version numbers associated with the domain group map cache for the site (2).

    - **Domain Group Cache WFE Update Result Set**, which returns the binary data needed to refresh the domain group map cache.

    - **Site Metadata Result Set**, which returns specialized site (2) metadata.

- **Event Receivers Result Set**, which returns information about the event receivers defined for the site (2).
- **Site Category Result Set**, which returns the categories of the site (2).
- **Site MetaInfo Result Set**, which returns the specialized site (2) metadata.
- **Site Feature List Result Set**, which returns a list of feature identifiers for the site collection that contains this site.
- **Site Feature List Result Set**, which returns a list of feature identifiers for the site (2).
- **Empty Result Set**, which is returned because the **METADATA_WEB** and **METADATA_WEB_NAVSTRUCT** flags are set and the site has no cached scope information.
- **List Metadata Result Set**, which returns metadata associated with the document list (1).
- **NULL Unique Permissions Result Set**, which is returned because the **METADATA_PREFETCH_SCOPES** flag is set, the **METADATA_LISTMETADATA_NOFETCH** flag has not been set, and permissions for the document list do not exist.
- **Event Receivers Result Set**, which returns information about the event receivers defined for the document list (1).
- **List Web Parts Result Set**, which returns information about the list (1) Web Parts defined for the document list (1).

36. The front-end web server builds a dynamic SQL query to retrieve information from the Sec_SiteGroupsView view. This query is sent to the SQL server using TDS.
37. The back-end database server returns a dynamic SQL result set containing information about the Team Site owners, the Team Site visitors, and the Team Site members (3).
38. The front-end web server builds a dynamic SQL query to retrieve information about the site (2) and document list (1) as it relates to the requesting user. This query is sent to the SQL server using TDS.
39. The back-end database server returns a dynamic SQL result set containing information about the current site (2) and document list (1).
40. The front-end web server builds a dynamic SQL query regarding what permissions the requesting user has within the site group. It does this by calling the proc_SecGetUserPermissionOnGroup stored procedure using TDS.
41. The back-end database server returns a dynamic SQL result set consisting of the output variables of the stored procedure.
42. The front-end web server builds a dynamic SQL query to retrieve information about the site (2) and document list (1) as it relates to the newly added user. This query is sent to the SQL server using TDS.
43. The back-end database server returns a dynamic SQL result set containing information about the current site (2) and document list (1).
44. The front-end web server redirects the end-user client to the "http://<YourSharePointServer>/_layouts/people.aspx?MembershipGroupId=5" page.

3.4 Example 4: Create a SharePoint Document Library File from the Client Console

This example describes a simple method for creating a file using the protocols covered in this system. This example uses the Creating a SharePoint Document Library File from the Client Console use case

described in section 2.5.1. This example is specifically for operations involving Microsoft SharePoint Foundation 2010.

Note The following steps and diagram consolidate multiple front-end web server to back-end database server actions, and multiple front-end web server to Active Directory actions into single flows. The step descriptions indicate where multiple actions are occurring and specify examples that provide more detail about those actions. In addition, the diagram and steps do not describe some initial interactions between the client and server that optionally happen on some clients, and which can also depend on whether the client has connected to the site previously to verify that the server is able to support WebDAV.

The main member protocols used in this sequence are [MS-WSSFO2] covering the stored procedures listed in the steps, and [MS-WDV].

The following assumptions are made for the purposes of this example:

- The user has read/write access permissions to an existing SharePoint Foundation 2010 document library called "http://server/site/doclib".
- The user is logged on to a client computer running Windows 7 operating system with an authenticated Windows session, and can access the Windows SharePoint Services site containing the document library.
- From a command prompt window on the end-user client machine, the user has typed the following command:

```
echo hello >\\server\site\doclib\hello.txt
```

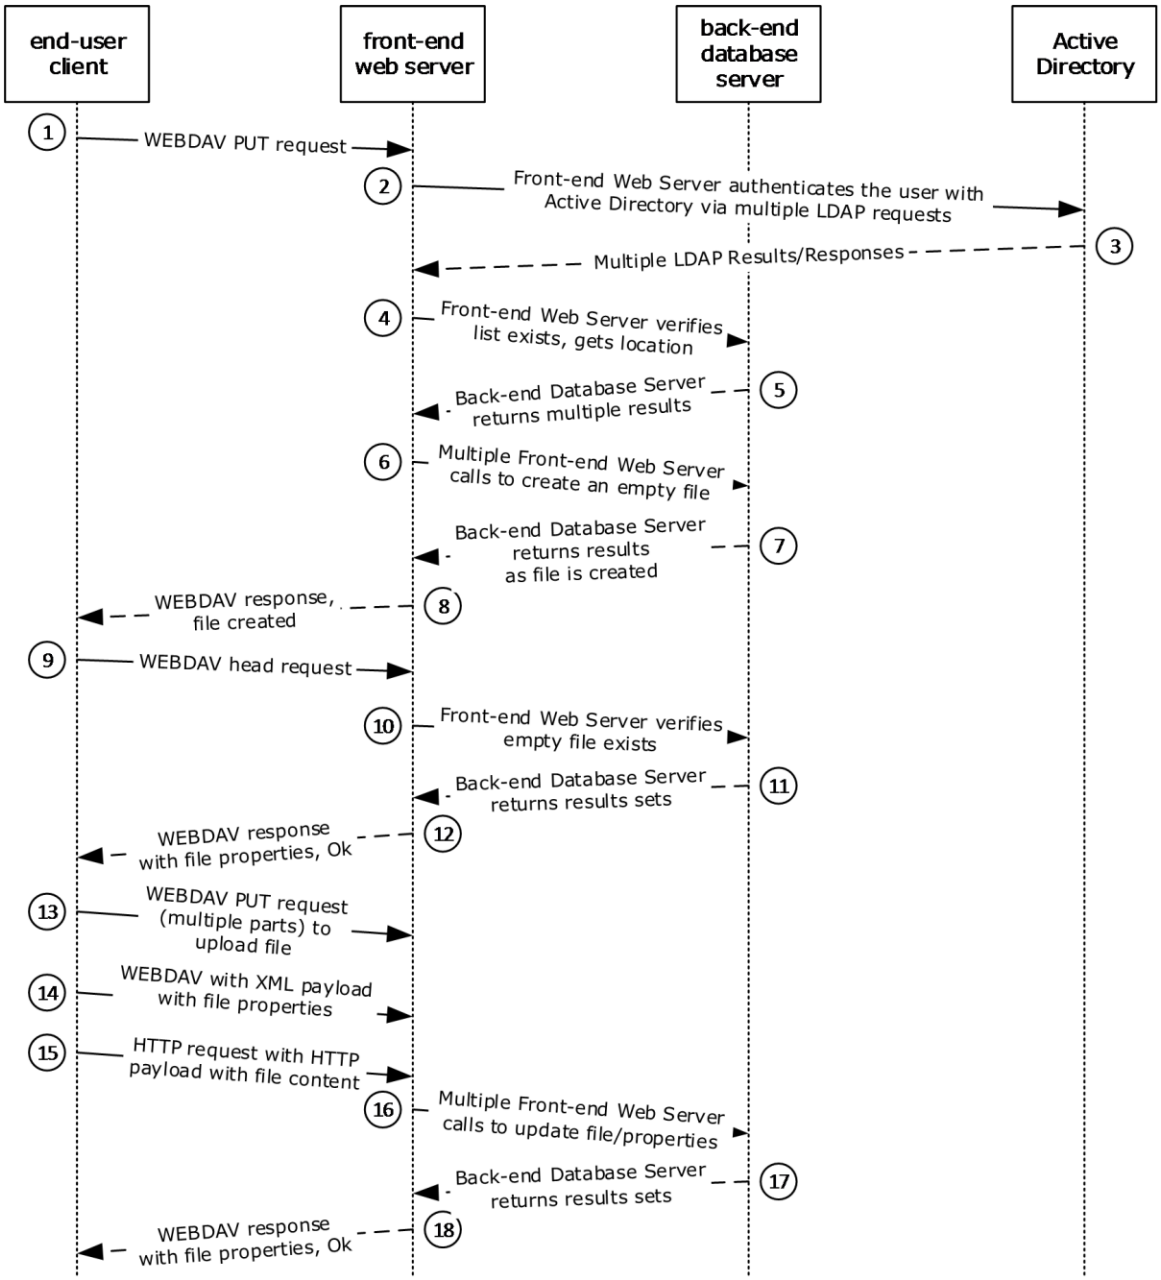

Figure 14: Create a SharePoint Document Library file from the client console

The [MS-WSSFO2] examples referenced for more details in the following steps use the SharePoint programming API; the actual steps can vary when the request is generated by user interaction with the front-end web server, as in this case.

The following actions happen:

1. The user initiates the `echo` command and the end-user client sends a WebDAV request to the server asking it to perform a **PUT** operation on the hello.txt file in the document library.

2. The front-end web server (IIS) authenticates the user with Active Directory. In practice, this can involve multiple LDAP requests with the Active Directory, especially if the user has not previously visited the site (2).

3. Active Directory responds with multiple LDAP results.

    For more information about authentication, see section 2.9.2.1.

    For more information about the scenario when the user has not previously visited the site, see [MS-WSSFO2] section 4.2 for SharePoint Foundation 2010.

4. In multiple roundtrips with the back-end database server, the front-end web server locates the content database for the document library, and verifies that the library exists.

5. The back-end database server returns multiple objects for the site collection, website and library to the front-end web server.

    For more information about steps 4 and 5, see [MS-WSSFO2] section 4.6 for SharePoint Foundation 2010.

6. In multiple roundtrips with the back-end database server, the front-end web server creates an empty file in the document library, and then, if successful, the front-end web server also verifies that the user has permissions to access and write to the document library.

7. The back-end database server returns multiple result sets as part of the process to create the file.

    For more information about file creation in steps 6 and 7, see [MS-WSSFO2] section 4.9 for SharePoint Foundation 2010.

8. The front-end web server returns a WebDAV response indicating the file was created successfully.

9. The client sends a WebDAV **HEAD** request to front-end web server with the URL to the hello.txt file in the document library, to verify the success of the previous call.

10. In multiple roundtrips with the back-end database server, the front-end web server retrieves the file.

11. The back-end database server returns multiple results sets as part of the process to retrieve the file.

    For more information about file retrieval in steps 10 and 11, see [MS-WSSFO2] section 4.1 for SharePoint Foundation 2010.

12. In response to the **HEAD** request, the front-end web server sends a response reporting that the request was successful.

13. Then the client sends a WebDAV **PUT** request to the front-end web server containing multiple parts, to upload the file and to update the file properties.

14. The client sends a WebDAV request to the front-end web server with an XML payload that has the file properties from the client.

15. The client sends a HTTP request to the front-end web server with an HTTP payload that has the file content; in this example, that content is simply the text "hello".

16. In multiple roundtrips with the back-end database server, the front-end web server updates the file and its properties in the document library.

17. The back-end database server returns multiple result sets as part of the process to update the files.

    For more information about file retrieval and update in Steps 16 and 17, see [MS-WSSFO2] sections 4.1 and 4.9 for SharePoint Foundation 2010. There is overlap with previous steps because stored procedures such as proc_FetchDocForUpdate ([MS-WSSFO2] section 3.1.5.21) are part of file updates as well as file creation.

18. Upon completing the update, the front-end web server sends a WebDAV response reporting that the request was successful.

4 Microsoft Implementations

- Windows Server 2003 operating system
- Windows Server 2008 operating system with Service Pack 2 (SP2)
- Windows Server 2008 R2 operating system

4.1 Product Behavior

<1> Section 2.5.1: Operating system versions other than Windows 7 operating system might require different steps than those specified in this use case.

5 Change Tracking

No table of changes is available. The document is either new or has had no changes since its last release.

6 Index

T

U

V