



Kevin Laahs
Emer McKenna
Veli-Matti Vanamo

Microsoft®

SharePoint

Technologies

Planning, Design
and Implementation

publishing
partner



Microsoft® SharePoint Technologies

Related Titles from

Digital Press

Kieran McCorry, *Microsoft Exchange 2003 Deployment and Migration*, ISBN 1-55558-316, 400pp, 2004

Pierre Bijaoui, *Scaling Microsoft Exchange Server*, ISBN 1-55558-239-7, 552pp, 2002

Sue Mosher, *Microsoft Outlook Programming: Jumpstart for Administrators, Developers, and Power Users*, ISBN 1-55558-286-9, 620pp, 2002

Jerry Cochran, *Mission-Critical Microsoft® Exchange 2003: Designing and Building Reliable Exchange Servers*, ISBN 1-55558-294-X, 480pp, 2003

Micky Ballardelli and Jan DeClercq, *Mission-Critical Active Directory: Architecting a Secure and Scalable Infrastructure*, ISBN 1-55558-240-0, 512pp, 2001

Mike Daugherty, *Monitoring and Managing Microsoft Exchange 2000 Server*, ISBN 1-55558-232-X, 432pp, 2001

Mike Daugherty, *Monitoring and Managing Microsoft Exchange 2003 Server*, ISBN 1-55558-302-4, 512pp, 2004

Kevin Laahs, Emer McKenna, and Don Vickers, *Microsoft SharePoint Portal Server: Building Knowledge Sharing Applications*, ISBN 1-55558-244-3, 544pp, 2002

Tony Redmond, *Tony Redmond's Microsoft® Exchange Server 2003 with SP1*, ISBN 1-55558-330-X, 1008pp, 2005.

Jan DeClercq, *Windows 2003 Server Security Infrastructures*, ISBN 1-55558-283-4, 800pp, 2004

Alain Lissour, *Understanding WMI: Exploiting Microsoft's Windows Management Instrumentation in Mission-Critical Computing Infrastructures*, ISBN 1-55558-266-4, 580pp, 2003

Alain Lissour, *Leveraging WMI: Using Windows Management Instrumentation to Solve Windows Management Problems*, ISBN 1-55558-299-0, 924pp, 2003

For more information or to order these and other Digital Press titles, please visit our Web site at www.books.elsevier.com/digitalpress

At www.books.elsevier.com/digitalpress you can:

- Join the Digital Press Email Service and have news about our books delivered right to your desktop
 - Read the latest news on titles
 - Sample chapters on featured titles for free
 - Question our expert authors and editors
- Download free software to accompany select texts

Microsoft® SharePoint Technologies

Planning, Design, and Implementation

Kevin Laahs
Emer McKenna
Veli-Matti Vanamo



ELSEVIER
DIGITAL
PRESS

Amsterdam • Boston • Heidelberg • London • New York • Oxford
Paris • San Diego • San Francisco • Singapore • Sydney • Tokyo

Elsevier Digital Press
30 Corporate Drive, Suite 400, Burlington, MA 01803, USA
Linacre House, Jordan Hill, Oxford OX2 8DP, UK

Copyright © 2005, Hewlett-Packard Development Company, L.P. All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher.

Permissions may be sought directly from Elsevier's Science & Technology Rights Department in Oxford, UK: phone: (+44) 1865 843830, fax: (+44) 1865 853333, e-mail: permissions@elsevier.com.uk. You may also complete your request on-line via the Elsevier homepage (<http://elsevier.com>), by selecting "Customer Support" and then "Obtaining Permissions."



Recognizing the importance of preserving what has been written, Elsevier prints its books on acid-free paper whenever possible.

Library of Congress Cataloging-in-Publication Data

Application submitted.

ISBN 1-55558-301-6

British Library Cataloguing-in-Publication Data

A catalogue record for this book is available from the British Library.

For information on all Elsevier Digital Press publications, visit our Web site at www.books.elsevier.com.

04 05 06 07 08 09 10 9 8 7 6 5 4 3 2 1

Printed in the United States of America

Contents

Foreword	xvii
Preface	xix
Acknowledgments	xxi
Part I Introduction to SharePoint Products and Technologies	
I How Did We Get Here?	3
1.1 From whence we came	3
1.1.1 Wasn't Exchange meant to be the collaboration platform?	3
1.1.2 What else was going on?	7
1.2 2003—A vintage year for collaboration	8
1.2.1 The collaboration ecosystem	9
1.3 Delivering the strategy	10
1.3.1 Exchange Server 2003	10
1.3.2 Windows SharePoint Services and SPS	11
1.3.3 Live Communications Server	12
2 Windows SharePoint Services	13
2.1 Providing the foundation for collaboration	13
2.2 Main architectural ingredients	14
2.3 A WSS site—major components and structures	16
2.4 Introducing Web Part Pages and Web Parts	19
2.5 Introducing site definitions, configurations, and custom templates	20
2.6 Lists	22
2.7 Views	26

2.8	Libraries	28
2.8.1	Document templates	28
2.8.2	Managing library items	29
2.9	Accessing WSS content	31
2.10	Actions	36
2.11	Alerts	36
2.12	A brief look at site management	36
2.13	Security	37
2.13.1	Site creation and user management	37
2.13.2	User rights	38
2.13.3	Site groups	38
2.13.4	Cross-site groups	39
2.13.5	Becoming a member of a site	40
2.13.6	Requesting access	42
2.13.7	Subsite permissions	43
2.13.8	Per-list permissions	43
2.13.9	Authentication	43
2.13.10	Anonymous access	44
2.13.11	Other considerations	45
2.14	SQL tables and WSS sites	45
2.14.1	SQL databases	46
2.14.2	Configuration database	46
2.14.3	Initial site creation	47
2.14.4	User information	50
3	“Office” Integration	53
3.1	Building on the framework	53
3.2	General WSS enhancements with Office	54
3.2.1	Uploading multiple items into a library	54
3.2.2	Creating new items in a library	56
3.2.3	Using the Outlook Address Book	58
3.3	Saving Office documents to WSS libraries	61
3.3.1	Property promotion	61
3.3.2	Opening and saving documents through Office applications	64
3.4	Working with WSS lists and Excel/Access	66
3.4.1	Importing or saving from Excel to a list	66
3.4.2	The Datasheet view	67
3.4.3	Synchronizing with Excel	69
3.5	The Shared Workspace task pane	70
3.6	Document workspaces	72

3.7	Meeting workspaces	76
3.8	Linking lists to Outlook	77
3.9	InfoPath and form libraries	80
3.10	Alerts and Outlook	82
3.11	Exchange public folders and WSS libraries	83
3.12	Outlook Web Access Web Parts	84
4	SharePoint Portal Server	87
4.1	Differentiating WSS and SPS	87
4.2	Major SPS components	88
4.3	What does SPS require of WSS?	89
4.4	Portal areas	91
4.5	A word about SPS URLs	95
4.6	Portal content navigation	97
4.7	Portal listings and news	99
4.8	Security	101
4.9	Site directory	102
4.10	Connecting team sites to a portal site	104
4.11	Alerts	105
4.12	Enterprise application integration	107
	4.12.1 Single sign-on	110
5	People and Personalization	113
5.1	Site personalization	114
5.2	Profiles	116
5.3	Personal sites (My Site)	119
5.4	Audience targeting	122
5.5	Smart Tag and Presence information	123
6	Search	127
6.1	Terminology	127
6.2	Major search components	128
6.3	Searching with WSS	131
6.4	Searching with SPS	136
6.5	Implementation and management options	139
	6.5.1 Content indexes	139
	6.5.2 Content sources	143
	6.5.3 Content source groups and search scopes	146
6.6	What does it all mean to the end user?	152

6.7	Combining WSS and SPS search	153
6.7.1	Two indexing engines = two lots of resources	156
6.8	Monitoring and customizing the search experience	157
6.8.1	Tweaking the Search Web Parts	157
6.8.2	Using the Search page: search.aspx	158
6.8.3	Calling the Search Web Service	158
6.8.4	Using the SPS Object Model	158
6.9	Shared Services considerations	159
6.9.1	Deploying a hierarchy of portal sites	159

Part II Planning, Deployment, and Operations

7 SharePoint Products and Technologies Architecture 167

7.1	Windows Server 2003 and IIS	167
7.1.1	Windows Server 2003 editions	167
7.1.2	Introduction to IIS 6.0	168
7.1.3	ASP.NET processing in IIS 6.0	171
7.1.4	Network load balancing	174
7.2	SharePoint Products and Technologies architecture	175
7.2.1	Building on top of IIS and ASP.NET	176
7.3	WSS architecture	178
7.3.1	Physical architecture components	180
7.4	SPS architecture	182
7.4.1	Physical architecture components	183
7.5	SPS Shared Services	187
7.6	Database services	188
7.6.1	Database communication	189
7.6.2	Database security options	190
7.7	Security architecture	191
7.7.1	Authentication and authorization	192
7.7.2	Data communication security options	194

8 Planning SharePoint Features 199

8.1	Getting to know the requirements	199
8.2	Solving your deployment jigsaw puzzle—which core server products might you need?	200
8.2.1	WSS or SPS 2003?	200
8.2.2	Microsoft Content Management Server 2002	202
8.2.3	Microsoft Office Project Server 2003	205
8.2.4	Microsoft BizTalk Server 2004	206
8.2.5	Live Communications Server 2003	207

8.3	Choosing the right account mode for WSS	208
8.4	Performing site-usage analysis	208
8.5	Retiring old sites with site-collection use confirmations	212
8.6	Setting site quota limits	215
8.7	Considering SPS Shared Services	217
8.8	Using SPS user profiles	219
8.9	Planning SPS search architecture	221
8.9.1	Which content types to index	222
8.9.2	Planning content indexes and content sources	223
8.9.3	Planning search scopes	225
8.10	Impact of Web Parts	226
8.10.1	Review of out-of-the-box functionality	226
8.10.2	Overview of custom Web Parts	227
8.10.3	Impact of Web Parts on the deployment	228
8.10.4	Managing and supporting custom Web Parts	232
8.11	Impact of customizations	233

9 Planning the Deployment Architecture 235

9.1	Overview of deployment sizing	235
9.1.1	Performance-testing methodology	236
9.1.2	Estimating required throughput for SharePoint Products and Technologies	237
9.1.3	Recommended application limits	239
9.2	Global deployment considerations	243
9.2.1	Addressing multiple languages	243
9.2.2	Distributed or regionalized server deployment	244
9.3	Sizing Windows SharePoint Services	246
9.3.1	Web front-end scalability	246
9.3.2	SQL server impact	247
9.3.3	Network impact	249
9.3.4	Choosing the right hardware topology	249
9.3.5	High-availability considerations	251
9.4	Sizing SPS 2003	252
9.4.1	Differences with Windows SharePoint Services	252
9.4.2	Choosing the right hardware topology	254
9.4.3	Search Service scalability	256
9.5	Disk subsystem planning	256
9.5.1	Database sizing	257
9.5.2	Search index sizing	258
9.6	Planning backup and restore	258
9.6.1	Out-of-the-box backup options	259

9.6.2	Backup-disk subsystem impact	260
9.6.3	Third-party backup solutions	261

10 Deploying SharePoint Technologies 263

10.1	Preparing for deployment	263
10.1.1	Preparing service accounts	263
10.1.2	Preparing the DNS architecture	266
10.1.3	Configuring the operating system for SharePoint deployment	268
10.1.4	Changing SQL Server ports	270
10.1.5	Configuring IIS for SharePoint Products and Technologies	271
10.2	Deploying Windows SharePoint Services	275
10.2.1	Installing Windows SharePoint Services	275
10.2.2	Configuring account and database settings	276
10.2.3	Configuring server settings and defaults	278
10.2.4	Extending virtual servers with Windows SharePoint Services	279
10.2.5	Server farm considerations	281
10.2.6	Configuring Windows SharePoint Services features	282
10.3	Deploying SharePoint Portal Server 2003	283
10.3.1	Installing SharePoint Portal Server 2003	283
10.3.2	Configuring account and database settings	287
10.3.3	Configuring SharePoint Portal Server topology	288
10.3.4	Server farm considerations	288
10.3.5	Creating SharePoint portal sites	290
10.3.6	Configuring Portal Server settings	291
10.4	Service packs and hotfixes	292

11 Migrating Data to SharePoint Products and Technologies 293

11.1	Migrating within SharePoint Products and Technologies	293
11.1.1	Tools and technologies to migrate between environments	294
11.1.2	Using Smigrate to move WSS spaces	296
11.1.3	Using STSADM to move content	298
11.1.4	When a company's structure changes	300
11.2	Approaching migrations from other platforms	302
11.2.1	Addressing functionality gaps and the user experience	303
11.2.2	Data structure	305

11.2.3	Versions and metadata	307
11.2.4	Identity mapping	308
11.3	Migration tools and technologies	309
11.3.1	Building a migration framework	309
11.3.2	Community effort: Custom SharePoint import/export tools	311
11.3.3	Commercial offerings	314
11.4	Migrating from legacy environments with the custom import/export tools	314
11.4.1	What data to move across	314
11.4.2	Migrating from Windows File Shares	316
11.4.3	Migrating from Exchange Public Folders	320
11.4.4	Migrating from Lotus QuickPlace	322
11.4.5	Migrating from Documentum eRoom	324
12	Supporting and Managing a SharePoint Deployment	327
12.1	Supporting SharePoint Products and Technologies	327
12.1.1	Planning a support strategy	327
12.1.2	Planning an operational model	330
12.1.3	Planning Service-Level Agreements	332
12.1.4	Expected support-call volume	334
12.1.5	Staffing an operational team	335
12.2	Proactive support strategy	336
12.2.1	Pathway to production	336
12.2.2	Active monitoring	337
12.2.3	Passive analyzing	339
12.3	Performance monitoring	339
12.3.1	Monitoring strategy	339
12.3.2	Identifying future bottlenecks	340
12.4	Managing SharePoint Products and Technologies	343
12.4.1	Managing sites and portal areas	343
12.4.2	Managing Search	344
12.4.3	Other common management tasks	344
12.4.4	Common support issues	345
Part III	Customizing and Extending SharePoint	
13	Introduction to SharePoint Development	349
13.1	Development architecture overview	349
13.1.1	IIS	352
13.1.2	Internet server API	352

13.1.3	ASP.NET handler	353
13.2	Enabling ASP.NET applications to coexist with SharePoint	355
13.3	SharePoint page rendering and Web Part overview	357
13.4	SharePoint site definitions and templates	359
13.5	SharePoint OM	361
13.6	SharePoint Web services	363
13.7	Developing SharePoint applications	363
14	Web Part Development	365
14.1	What is a Web Part?	365
14.2	Creating a Web Part Page	367
14.3	Adding, modifying, and arranging Web Parts	368
14.3.1	Adding Web Parts	369
14.3.2	Modifying Web Parts	373
14.3.3	Arranging Web Parts	378
14.3.4	Removing a Web Part from a Web Part Page	380
14.4	Creating a simple Web Part	380
14.4.1	Web Part component files	381
14.4.2	Getting started	382
14.4.3	Coding the Web Parts	386
14.4.4	Deploying the Web Part	391
14.5	Web Part infrastructure	392
14.5.1	Web Part rendering process	392
14.6	Web Part properties and methods	394
14.6.1	Implementing child controls	397
14.6.2	Customizing the Web Part menu	400
14.6.3	Tool Panes and Tool Parts	404
14.7	Web Part connections	417
14.7.1	Connection interfaces	418
14.7.2	Connection processing	420
15	Deploying and Debugging Web Parts	437
15.1	Code Access Security	437
15.1.1	Trust level	438
15.1.2	Policy files	439
15.2	CAS implementation	446
15.2.1	Raising the trust level for all SharePoint sites	447
15.2.2	Deploying Web Parts to the GAC	447
15.2.3	Create a custom policy file	448

15.3	Deploying Web Parts	451
15.3.1	Packaging Web Parts	453
15.3.2	Creating a cabinet file	454
15.3.3	Deploying a Web Part Package	456
15.3.4	Web Part Galleries	460
15.4	Debugging Web Parts	462
15.4.1	Configuring SharePoint Server	464
15.4.2	Configuring a remote computer	465
16	Customizing SharePoint Using FrontPage 2003	467
16.1	Microsoft Office FrontPage 2003	467
16.2	Editing a SharePoint site with FrontPage 2003	467
16.3	Features of FrontPage 2003	468
16.3.1	WYSIWYG editing	469
16.3.2	Quick Tag tools	470
16.3.3	Image tracing	472
16.3.4	Table and cell layout	473
16.3.5	Themes	474
16.3.6	Dynamic templates	476
16.4	Web Part and list integration	485
16.4.1	Web Part Pages	486
16.4.2	Web Part zones	487
16.4.3	Web Parts	488
16.4.4	Lists and libraries	494
16.5	Data views	498
16.5.1	Data sources and Data views	499
16.5.2	The Data Source Catalog	499
16.5.3	Inserting the Data view	504
16.5.4	Selecting the table style	505
16.5.5	View style options	508
16.5.6	Adding fields to the Data view	508
16.5.7	Formatting the Data view	511
16.5.8	Filtering data	511
16.5.9	Sorting and grouping data	513
16.5.10	Conditional formatting	515
16.5.11	Web Part formatting	516
16.5.12	Inserting subviews	517
16.5.13	Connecting Data views	518
16.5.14	Converting lists to Data views	519

16.6	Customizing menus	520
16.6.1	Customizing the navigational structure	520
16.6.2	Collapsing navigation bars	523
16.7	Web Packages	528
16.8	Web Log Web Packages	529
16.8.1	What is blogging?	530
16.8.2	Functionality overview	531
16.8.3	Creating a Web log site	541
16.8.4	Setting up your Web log site	542
16.9	Ghosting and FrontPage 2003	545
16.10	Setting up your site to restrict access through FrontPage	547
16.11	FrontPage 2003 features disabled when editing an SPS site	548

17 Customizing and Extending SharePoint 551

17.1	Basic customizations	551
17.1.1	Changing the title of the Web Part Page	553
17.1.2	Replacing the logo	554
17.1.3	Merging the SPS horizontal menu bars	556
17.2	Customizing the SharePoint stylesheets	558
17.2.1	Creating the custom stylesheet	560
17.2.2	Removing the logo that appears on the menu bar	561
17.2.3	Changing the background color of the top menu bar	562
17.2.4	Reducing the height of the menu bar	562
17.2.5	Removing the background color associated with the search boxes	564
17.3	Custom themes	565
17.4	Site definitions	568
17.4.1	ONET.XML	570
17.4.2	Site-definition banners and footers	591
17.4.3	WEBTEMP.XML	592
17.5	Creating a custom site definition	595
17.6	Custom templates	596
17.6.1	Site definition advantages	600
17.6.2	Site definition disadvantages	600
17.6.3	Custom template advantages	600
17.6.4	Custom template disadvantages	600

17.7	SharePoint object model overview	601
17.7.1	Creating a SharePoint ASP.NET Web application	603
17.7.2	Coding a SharePoint application	605
17.8	Leveraging the event architecture	609
17.9	Accessing SharePoint data using Web Services	620
17.10	Conclusion	628

Index**629**

This Page Intentionally Left Blank

Foreword

I love software that runs and delivers value out of the box. SharePoint Portal Server (SPS) 2001 certainly delivered great value to customers, because it was easy to deploy and use. SPS 2001 also affected the economics of the portal market, because its pricing made it viable for companies to dip their toes into portal deployment without incurring huge cost. All software evolves over time, and we now have a new version of SPS, plus the addition of Windows SharePoint Services as part of the base Windows operating system. So, now is a good time to take the opportunity to consider the value that companies can extract from SharePoint.

While SharePoint delivers value out of the box, you'll only get limited value if you accept everything as is. For the software to be truly valuable, I believe that you have to consider the needs of the business and then look at how SharePoint can help you address those needs. Bending SharePoint to meet what you think the business wants is probably pointless. Sometimes you'll succeed, but most of the time you'll end up with frustrated and unhappy users. Experience is tremendously important, too, because without experience, you'll never quite know where the limits exist and where you can apply tweaks to satisfy business needs.

This book is late. I have moaned at the authors on this point for a long time, but they quote my own arguments back to me, citing the need to capture their experience of meeting real business needs within HP as we deploy and use SPS and Windows SharePoint Services to help our people be more productive. Experience takes time to accumulate, so I guess I shouldn't be surprised that the book is late. Now that it has appeared, the content in the book shows real experience of technology in action in many places, and that's its real value. Many people can sit down, learn a technology, and describe it in a book. Few can tell the story from the perspective of experience and make their hard-earned knowledge available in an accessible form. Although it took time, I think this book hits the target, and I'm sure it will

be valuable to you, especially if you face the challenge of migrating from the original SharePoint architecture as implemented in the 2001 release to the new technology platform exploited by SPS 2003. Migrating data is easy enough; the challenge exists in understanding how to best exploit the new architectural capabilities and what you need to do with the customized Web Parts and other tweaks that administrators and programmers deploy to meet business requirements.

As with any book, you have to take the information presented here and put it into context with your own experience, requirements, and pressures (time and other). Dip into the book, find its delights, and use them to your advantage.

Enjoy!

Tony Redmond
Vice President and Chief Technology Officer
HP Services

Preface

Mergers, mergers, mergers . . .

SharePoint Products and Technologies are the second generation of the SharePoint family, whose roots are to be found in the Version 1 products called SharePoint Team Services and SharePoint Portal Server (SPS) 2001. Since our first book on SPS 2001, we have lived through many forms of mergers (both personal and business, including the largest IT merger in history between HP and Compaq) and have learned that great things can come from successful mergers. Successful mergers tend to leverage the individual strengths of all involved into a single, stronger force that benefits all concerned.

And this is also the case with SharePoint. The very loosely coupled Version 1 products have come together to produce a more tightly integrated and richer platform, upon which successful collaboration and information sharing can occur. As well as finally delivering on the SharePoint brand, the overall offering merges nicely with other major Microsoft directions, such as Windows 2003, the .NET Framework, and Visual Studio .NET. The delivery of the SharePoint platform also offers benefits to others that play in the collaboration space, as seen with the likes of Office, Project, and InfoPath as they consume underlying SharePoint services.

We, the authors of this book, all work for HP Services and are involved with SharePoint at many different levels—through consulting directly with customers, presenting at industry conferences on SharePoint technologies, contributing to industry publications, and architecting and building HP Service’s internal knowledge network based upon SharePoint. We hope that the range of our experience helps you achieve success in your encounters with SharePoint.

Book structure

This book does not cover every single aspect of SharePoint, and if we have missed your favorite topic, then we apologize. We do feel, however, that we have tackled the most important areas, as well as those subjects in which our experience complements the standard product documentation and other books that are out there. We urge you to use your favorite search engine if the answer you seek is not contained here!

We understand that your interest in SharePoint may derive from many different sources, and therefore, we have structured the book into three distinct parts. Part I is for those of you who want to find out about what SharePoint can actually do for you in end-user terms. If you are new to SharePoint and want to see if it can do a job for you, then we recommend you read Part I.

Part II is for system architects who want to know how to plan, design, implement, and manage a SharePoint environment. If you are tasked with rolling out SharePoint to support your business, then we recommend that you read Part II.

Part III is for those who want to extend base functionality to meet specific business needs. If you are a developer who is tasked with enhancing a SharePoint environment, then we recommend that you read Part III.

Note that there will be some slight repetition among the three parts in places. We are aware of this, and it was a deliberate choice in order that you could read each part in isolation from the others. The repetition is merely to ensure that you have the prerequisite information to follow along with the current topic.

Where there's a merger, there's a divestiture . . .

Almost as common as mergers are divestitures, and again, we have experienced many over the past years. We're glad to say that none of us has experienced divorce, albeit writing a book could be cited as a justifiable reason in any such action. We just wanted to say that Don Vickers, who coauthored the first book, could not join us for this second journey and we wish him continued success in all that he does. We did manage to find an excellent replacement in Veli-Matti Vanamo, who, aside from being an alphabetically handy replacement, has also matured sufficiently to appreciate (and spell correctly) Laphroaig.

Kevin, Emer, Veli

Acknowledgments

It is with great relief that we find ourselves writing this section of the book; there were times during this process that we didn't think we'd make it, but here we are! While writing a book provides many benefits, including an opportunity for personal growth, it requires an enormous commitment—not only from the authors but from many other sources. The acknowledgments section gives us each a chance to express our thanks to the people who helped to get us here by providing their ongoing advice, encouragement, and support.

We'd like to thank the publishing team—namely Alan Rose, Lauralee Reinke, and Theron Shreve—for their incredible patience, dedication, flexibility, commitment, and support throughout. You truly are a wonderful team and it was a pleasure working with you. Special thanks go to Tony Redmond for his support and providing the Foreword; Mike Topalovich, Gary Adams, and Adam Rosenblatt for taking the time to provide their reviews; and Eric Tipton for his contribution to the review process.

We are fortunate to work inside a company that has a wealth of experience and a sharing culture that allows us to draw on this experience in many ways. So thanks to all our colleagues at HP. We are also fortunate to have a very close working relationship with Microsoft, and we'd like to thank all our friends there for their willingness to help us and answer our many questions—we promise to continue bombarding you guys with e-mails!

Before we move on to our individual acknowledgements, collectively we'd like to thank those of you who ordered this book many months ago and didn't cancel your order—thank you so much for your patience; we hope you find it worth the wait!

Personal Acknowledgments

To Emer and Veli, my family and friends, and to my wife, Wendy, and children, Jenny and Euan—thanks for putting up with me.

—Kevin Laahs

I'd like to thank Kevin and Veli for their knowledge, expertise, and dedication—as always, it was great working with you guys.

Most of all, I would like to thank my wonderful husband, Michael, and my beautiful little girl, Caoimhe, for supporting my contribution to this book, which required countless nights and weekends away from both. I am so relieved that we can spend time as a family once more. I love you both dearly. XOXO

—Emer McKenna

First of all, I would like to thank Kevin and Emer for the opportunity to co-author the book, and for the countless hours you folks spent pulling me through it. I'd like to thank my wife, Audrey, for her enormous support and even agreeing to go through with the wedding during the lengthy writing process. Finally, many thanks to my family for all the advice and support throughout the years.

—Veli-Matti Vanamo

Part I

Introduction to SharePoint Products and Technologies

This Page Intentionally Left Blank



How Did We Get Here?

1.1 From whence we came

This is our second book with the word “SharePoint” in its title. Our first, entitled *Microsoft SharePoint Portal Server: Building Knowledge-Sharing Applications* (2001), dealt primarily with Microsoft SharePoint Portal Server 2001 (SPS 2001). I guess it would be natural to assume that this offering is merely an update covering the next release of this product—Microsoft Office SharePoint Portal Server 2003 (SPS 2003)—but this is a false assumption. This book actually covers Microsoft SharePoint Products and Technologies, an umbrella name for those offerings from Microsoft that facilitate collaboration. Now, while SPS is still a major player in this area, it is only part of a strategy designed to connect people to information across all levels of an organization.

This strategy is significantly different from the one that was around when we wrote our first book. It is therefore important to understand how Microsoft has arrived at such a strategy, as this will help you appreciate the thinking behind the current product offerings. Hopefully, this will then assist you in making informed decisions about how to deploy a SharePoint environment for your specific collaborative needs.

1.1.1 Wasn’t Exchange meant to be the collaboration platform?

If you cast your mind back to 2000, you will recall the launch of Exchange Server 2000. A significantly different architecture from that of Exchange 5.5, it was hailed as not just a messaging server, but also as a collaboration server. Additionally, it brought new technologies to the table that would be used by other Microsoft product offerings. Witness here the Web Storage System, which is at the core of SPS 2001.

As a collaboration server Microsoft was clearly trying to offer something that could compete with Lotus Notes. Few would argue that Exchange provided the best messaging platform, while Notes provided the best collaborative application platform, and both products were trying to make ground in their “weaker” areas. Thus, a major focus of Exchange 2000 was providing a platform upon which developers could build collaborative applications in addition to delivering the best-in-class messaging server. As it turned out, Exchange’s strength as a messaging server alone was enough to keep the threat of Notes at bay, and this is one reason why Microsoft has subsequently de-emphasized Exchange as a collaboration platform and focused it firmly as a pure messaging server. Indeed, you only have to look at Exchange Server 2003 and note that real-time collaboration options such as Instant Messaging, which were part of Exchange Server 2000, are no longer part of the Exchange offering.

However, the consideration that the messaging battle had been “won” was not the only reason for the de-emphasis. The failure of Exchange Server 2000 to become a ubiquitous collaboration platform was another. True, some applications were built on top of Exchange, but it basically failed to capture the interest of the majority of application developers out there. It also met with resistance from administrators, whose primary purpose in life was to ensure a highly performing messaging server. E-mail is mission critical for most organizations these days, and anything that might jeopardize its availability or service-level agreements is met with scorn!

Understanding the reasons for the failure has helped shape the current strategy. At the time, Microsoft had three major goals for Exchange as a collaborative application development platform:

1. Make Exchange-based data widely available through many protocols and APIs. If the Web Storage System were to be the place to store all sorts of information, then it would have to be ubiquitously available.
2. Provide an extensible platform. In order to meet a wide variety of needs, new services would need to integrate easily into the platform.
3. Deliver a compelling environment for the development of rich, collaborative applications. Encourage developers to use the Web Storage System as their platform of choice.

Let’s take a look at each goal and see how it fared. First, Microsoft certainly did provide ubiquitous access to Exchange data. The primary APIs were ActiveX Data Objects (ADO), Collaboration Data Objects for

Exchange (CDOEX), and Web-based Distributed Authoring and Versioning (WebDAV). The support of ADO meant that developers who were used to coding against, say, SQL or Oracle, could code against the Web Storage System in much the same way. This met with Microsoft's goal of making the Web Storage System appeal to the masses. CDOEX provided a rich API that made it easy to handle typical Exchange items such as messages, calendars, contacts, and tasks. WebDAV support made building Web-based clients using the HTTP protocol a reality—Outlook Web Access (OWA) being a prime example here. Indeed, using any mixture of these three APIs, you could quite easily access and manipulate items in user mailboxes and public folders.

Sounds like a pretty good picture for data access, so why wasn't it successful? Well, while ADO and CDOEX were robust and powerful APIs, they could only run on the Exchange Server. This meant that applications themselves would have to be installed and run on the Exchange Server. "Not on MY Messaging Server!" would be the response of most Exchange administrators. You see, administrators tend to be passionate about their messaging service and do not want to risk compromising the service it provides to end users. Therefore, it was very difficult to put forward a good case for installing such applications onto the Exchange Server. The only other alternative would be to supply Exchange Servers that were separate from the core messaging service specifically for these applications—clearly a prohibitive expense!

WebDAV is different from ADO and CDOEX in that it can run remotely. In other words, you can run your custom application on one server and have it access Exchange-based data on another server. Now, while this approach appeased Exchange administrators, writing raw WebDAV applications was rather complex. Thus, developers did not embrace this as a viable solution.

Finally, with the provision of these multiple APIs, there were multiple ways to achieve the same aim. For example, there are eight different ways to create a folder in the Web Storage System, so which is the right one to use? This plethora of options merely confused rather than assisted the developer.

Providing an extensible platform was achieved mainly through the support of events at three different levels—SMTP Protocol, SMTP Transport, and Store. Events allowed you to write your own custom code that could react appropriately as messages traversed your environment and as items were manipulated in the Web Storage System. A couple of examples will help clarify the power of events. First, SPS 2001 actually relied heavily on a Store event in order to ensure the integrity of documents as they were

checked in and published (remember SPS 2001 uses the Web Storage System that was initially provided by the Exchange group). As documents are created in an SPS 2001 document library, the Store event ensures that valid profiles are filled in, version numbers are maintained, and so forth.

The second example is that of third-party antivirus vendors, who typically harness the SMTP Transport events in order to scan messages before they are actually delivered to a user's mailbox. This is a great example of how new services could be layered easily on top of the platform.

Now, much like the problems with ADO and CDOEX, any application utilizing events would ultimately have to be installed directly onto the Exchange Server. While it is true that SMTP events could run on boundary servers not hosting user mailboxes, Store events would typically have to be on mailbox and public folder servers if they were to add any significant value. Yet, again, Exchange administrators hollered, "Not on MY Exchange Server!" A secondary issue that relates particularly to SMTP events was that you really needed to be able to code in a language such as C++ in order to exploit their capabilities fully. This alienated a lot of those people who would not consider themselves to be hard-core developers and meant that the feature was not overly utilized, as it did not appeal to the masses.

The third goal of delivering a compelling platform for the development of collaborative applications had its challenges as well. Microsoft wanted to lure developers into using the Web Storage System as their primary platform, not just for the data associated with an application but for all its collateral, including source code, forms, and the like. As such Microsoft envisaged all application code and data to be stored within public folders. To this end it provided supporting technologies and showcase applications such as Web Storage System Forms, FrontPage Extensions for Web Storage System Forms, sample public folder-based applications, a workflow engine/designer, and regular updates to the Exchange Software Developer Kit. Interestingly, the workflow engine actually leveraged Store events and was a great example of how the overall platform could be extended.

Did these features attract the application developer? Unfortunately, the answer was no, due mainly to the fact that the whole paradigm was confusing and inconsistent when compared with other Microsoft initiatives. For example, writing a Web application with Visual Studio 6.0 was fairly straightforward for a developer, but there was no direct support for utilizing the Web Storage System as the storage mechanism for such an application. Furthermore, the whole approach seemed detached from the main development story emanating from Microsoft—that being the .NET Framework, Visual Studio .NET, and the whole Web Services world.

So the bottom line here was that the strategy was confusing, inconsistent, and impractical, and there was a big message from administrators saying, “Don’t mess with my messaging server!” So, clearly, another approach had to be taken to find a suitable platform for building applications that would enhance the lives of information workers in the collaboration area.

Note that the effort made to make Exchange a collaboration platform was not all in vain. I say this because there are pieces of the strategy that align with the future and have made Exchange a good citizen in a .NET world—that being the support of WebDAV such that Exchange services can be consumed by remote applications as and when required. As we will see later, this capability is actually harnessed by Windows SharePoint Services in order to allow documents to be e-mailed into a document library.

1.1.2 What else was going on?

Ever heard of a product called SharePoint Team Services (STS)? I’ll bet you probably have, but at the time of its release, SPS 2001 was getting all the airplay. In fact, in our first book we dedicated a whole paragraph to STS (on page 19 if you are really interested), which just goes to show that we did not really understand its true purpose in life (hey, at least we are honest!).

In our defense, Microsoft did not actually give STS the exposure it deserved at that time, and its relation to SPS was purely nominal. Part of the Office XP family, STS came free with a FrontPage 2002 license and, from a technical point of view, had its technology roots firmly planted in FrontPage Server Extensions. It was designed to allow the easy creation of Web sites that users could use as a mechanism for sharing information and collaborating.

It is fair to say that Microsoft did not market the capabilities of STS as well as it could have and put more effort into the technically unrelated SPS product that had shipped a few months previously. The confusion over these two products initially caused STS to lose momentum and not attain its potential in terms of customer uptake. In fact, a common misunderstanding was that an STS team site could only handle groups of up to 50 users. Only after Microsoft publicized its internal usage did STS gain credibility as concerns over scalability were alleviated. With additions such as Self Service Creation, STS proved to be a very popular platform for empowering users with the ability to create team sites instantly and start collaborating without involving IT groups for site creation.

SPS and STS both had their strengths, but were architecturally different and used different development metaphors. For example, to build on SPS

required knowledge of the Digital DashBoard Architecture and Web Parts, whereas building on STS required skills with FrontPage Remote Procedure Protocol (RPC) and Collaborative Markup Language (CAML).

However, STS filled a crucial role that couldn't be fulfilled by SPS 2001—it provided an environment in which ad hoc teams could come together, work on some collateral, and then disband. For example, a sales team could get together to work on a new proposal for a customer, or three people could cooperate on the creation of a book. SPS 2001 was designed more to support the publishing and discovery of information than the creation of such information and, therefore, did not lend itself well to the needs of team-based collaboration. The lesson learned is that an effective knowledge-sharing environment must support both the dynamic creation and publishing of information.

Collaboration means working together, especially in a joint intellectual effort. It's fair to say that although the three main offerings from Microsoft in the collaboration space each had its individual merits (Exchange Server 2000, SPS 2001, and STS), they did not exactly benefit from collaborative engineering. Clearly, a more cohesive, consistent, and all-encompassing message was required, and, indeed, this is where the much more compelling story of SharePoint Products and Technologies enters the fray. Which brings us to 2003.

I.2 2003—A vintage year for collaboration

Microsoft released a bundle of products in 2003. This is not really surprising, as it releases many products each year, but “working together” and “joint intellectual effort” are clearly emerging product trends.

The Microsoft Office family of products has been significantly extended, with members sharing much more than just the name. The full list of Office 2003 products is as follows:

- Access
 - Excel
 - FrontPage
 - InfoPath
 - Live Communications Server
 - Live Meeting
 - OneNote
-

- PowerPoint
- Project
- Publisher
- SharePoint Portal Server
- Visio
- Word

The interesting point here is that formerly “stand-alone” server products have now been brought into the family—in particular SPS and Live Communications Server. This indicates a desire to extend the functionality of personal productivity applications through server-based services.

Completing the major 2003 product set is Windows Server 2003, Windows SharePoint Services, and Microsoft Exchange Server 2003. Note that, as discussed earlier, Exchange 2003 has dropped the “Collaboration” tag and is now firmly positioned as a messaging server. Based on the lessons learned with Exchange 2000, SPS 2001, and STS, the collaboration strategy now focuses on an extensible infrastructure supporting services that citizens of multiple collaborative environments can consume, rather than on one that assumes people only need to collaborate from within their messaging system.

The basis of the collaboration story is now rooted inside Windows itself, through Windows SharePoint Services, and through individual product offerings that can easily leverage these services.

1.2.1 The collaboration ecosystem

As people go about their day-to-day jobs, they encounter many opportunities for collaboration. Indeed, the need to work together with one or more people (or with collateral created by other people) can arise at any moment and is seldom restricted solely to when you are working within a specific application. For example, as I author this part of our book, I think of questions that can be answered by my colleagues, I search for other collateral in the same topic area, and I think of tasks that need to be completed before this book can be published. People collaborate on many different levels and from within many environments. Looking at a pre-2003 Microsoft-centric environment, we see that one-to-one communication is performed via Outlook and Instant Messaging; teamwork on ad hoc projects is performed through a mixture of file shares, public folders, and SharePoint Team Services v1.0; conferencing is performed through either NetMeeting

or Exchange Conferencing Server, and publishing is performed through SPS 2001.

Microsoft's collaboration strategy is to provide services that can be consumed "naturally" as information workers go about their day-to-day jobs. By this, I mean that the information workers should not have to divert from their current focus when they need to collaborate. If I am authoring a document in Word and something in the document triggers the need to allocate a task to a colleague, I should be able to do so directly from Word itself and not have to go searching for another application. Collaborative services should be baked into the operating environment and called upon as and when it makes sense.

Another part of the strategy builds upon the realization that in order to collaborate effectively on a piece of work, users need to access multiple information sources—and not just document sources. Connecting people and information is very important. For example, if I am researching a topic and come across an interesting piece of work by a particular person, I may want to see other items that have been authored by that same person or may want to converse immediately with that person through an Instant Messaging conversation. Alternatively, I may want to identify people that have a particular skill set and target them for specific questions. Connecting people and information is therefore something addressed by the strategy.

The strategy therefore is to deliver rich collaboration services built into the infrastructure that go far beyond that of traditional file-sharing services, but are just as easy to use. Endusers are empowered to use these services through the tools they most commonly use—Office, Windows Explorer, and Internet Explorer—and other applications can avail themselves of and build on top of these infrastructure services.

1.3 Delivering the strategy

Four server products form the backbone of the collaboration strategy (Figure 1.1).

1.3.1 Exchange Server 2003

Focusing on Exchange Server 2003's strength as a messaging server, we see that previous collaboration services have been removed from the product. Exchange no longer delivers Instant Messaging and the Exchange Conferencing Service. Exchange also provides the best platform for personal information management in terms of calendars, contacts, and tasks and enables

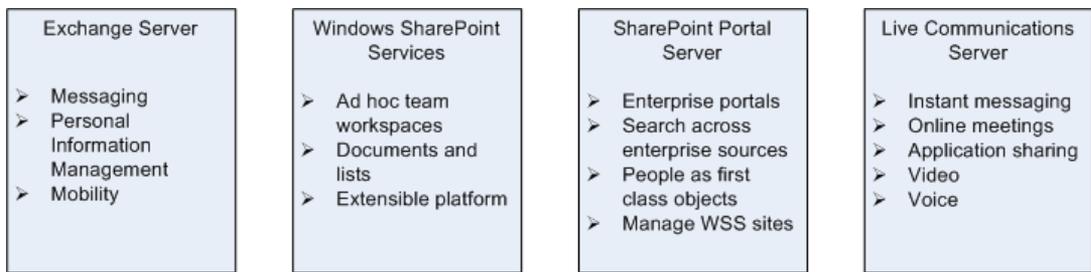


Figure 1.1 *Major products in the collaboration space.*

you to access your mailbox through multiple mobile clients, including Outlook Web Access and Outlook Mobile Access.

I.3.2 Windows SharePoint Services and SPS

Although these products could be seen as upgrades to SharePoint Team Services v1.0 and SPS 2001, it is better to think of them as new products, delivering the core part of Microsoft's collaboration strategy. Together they are known as SharePoint Products and Technologies, and Microsoft had several goals when designing the SharePoint suite.

First, it wanted to provide a scalable base service for collaboration that could be leveraged by many applications. Second, it wanted to ensure that information sharing could happen at all levels of an organization—groups of people, teams of people, departments, and organizations. Third, it wanted to ensure that this environment exploited the .NET Framework and was easily consumed and extended by third-party applications. Finally, it wanted to deliver on the SharePoint brand by ensuring that any SharePoint product utilized the underlying SharePoint services.

Windows SharePoint Services (WSS) is so called because it is part of Windows (available as a free download from Microsoft's Web site), and it provides the base collaborative infrastructure, as we will see in Chapter 2. It is an extensible Web-based environment that lets users quickly provision dynamic Web sites for specific purposes. For example, users can provision a Web site in order to host the collateral of a project or to collaborate on the creation of a particular document. The services offered by WSS are consumable via an object model and Web Services and can therefore be used in any application. Refer to Chapter 3 for details of how Office 2003 consumes the services provided by WSS. WSS provides a scale-out and scale-up three-tier architecture utilizing ASP.NET to deliver Web pages and SQL Server to host data.