

THE EXPERT'S VOICE® IN SQL SERVER

Pro SQL Server Wait Statistics

*A PRACTICAL GUIDE TO ANALYZING
PERFORMANCE IN SQL SERVER*

Enrico van de Laar

EXTRAS ONLINE

Apress®

www.allitebooks.com

Pro SQL Server Wait Statistics



Enrico van de Laar

Apress®

Pro SQL Server Wait Statistics

Enrico van de Laar
De Nijverheid
Drachten, The Netherlands

ISBN-13 (pbk): 978-1-4842-1140-3
DOI 10.1007/978-1-4842-1139-7

ISBN-13 (electronic): 978-1-4842-1139-7

Library of Congress Control Number: 2015950153

Copyright © 2015 by Enrico van de Laar

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law.

Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, logo, or image we use the names, logos, and images only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Managing Director: Welmoed Spahr

Lead Editor: Jonathan Gennick

Technical Reviewer: Chris Presley

Editorial Board: Steve Anglin, Mark Beckner, Gary Cornell, Louise Corrigan, Jim DeWolf,

Jonathan Gennick, Robert Hutchinson, Michelle Lowman, James Markham,

Susan McDermott, Matthew Moodie, Jeffrey Pepper, Douglas Pundick, Ben Renow-Clarke,

Gwenan Spearing, Matt Wade, Steve Weiss

Coordinating Editor: Jill Balzano

Copy Editor: April Rondeau

Compositor: SPi Global

Indexer: SPi Global

Artist: SPi Global

Cover Designer: Anna Ishchenko

Distributed to the book trade worldwide by Springer Science+Business Media New York, 233 Spring Street, 6th Floor, New York, NY 10013. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail orders-ny@springer-sbm.com, or visit www.springeronline.com. Apress Media, LLC is a California LLC and the sole member (owner) is Springer Science + Business Media Finance Inc (SSBM Finance Inc). SSBM Finance Inc is a **Delaware** corporation.

For information on translations, please e-mail rights@apress.com, or visit www.apress.com.

Apress and friends of ED books may be purchased in bulk for academic, corporate, or promotional use. eBook versions and licenses are also available for most titles. For more information, reference our Special Bulk Sales–eBook Licensing web page at www.apress.com/bulk-sales.

Any source code or other supplementary material referenced by the author in this text is available to readers at www.apress.com. For detailed information about how to locate your book's source code, go to www.apress.com/source-code/. Readers can also access source code at SpringerLink in the Supplementary Material section for each chapter.

Printed on acid-free paper

Contents at a Glance

About the Author	xvii
About the Technical Reviewer	xix
Acknowledgments	xxi
Introduction	xxiii
■ Part I: Foundations of Wait Statistics Analysis.....	1
■ Chapter 1: Wait Statistics Internals.....	3
■ Chapter 2: Querying SQL Server Wait Statistics	21
■ Chapter 3: Building a Solid Baseline	47
■ Part II: Wait Types.....	69
■ Chapter 4: CPU-Related Wait Types	71
■ Chapter 5: IO-Related Wait Types	97
■ Chapter 6: Backup-Related Wait Types	135
■ Chapter 7: Lock-Related Wait Types	145
■ Chapter 8: Latch-Related Wait Types.....	171
■ Chapter 9: High-Availability and Disaster-Recovery Wait Types.....	207
■ Chapter 10: Preemptive Wait Types.....	227
■ Chapter 11: Background and Miscellaneous Wait Types.....	247

■ CONTENTS AT A GLANCE

- **Chapter 12: In-Memory OLTP–Related Wait Types..... 269**
- **Appendix I: Example SQL Server Machine Configurations 285**
- **Appendix II: Spinlocks..... 289**
- **Appendix III: Latch Classes 293**

- Index..... 301**

Contents

About the Author	xvii
About the Technical Reviewer	xix
Acknowledgments	xxi
Introduction	xxiii
■ Part I: Foundations of Wait Statistics Analysis.....	1
■ Chapter 1: Wait Statistics Internals.....	3
A Brief History of Wait Statistics	3
The SQLOS.....	5
Schedulers, Tasks, and Worker Threads.....	8
Sessions	9
Requests.....	10
Tasks.....	10
Worker Threads	11
Schedulers.....	14
Putting It All Together.....	15
Wait Statistics	16
Summary.....	19

- **Chapter 2: Querying SQL Server Wait Statistics 21**
 - Sys.dm_os_wait_stats 21
 - Sys.dm_os_waiting_tasks 23
 - Querying sys.dm_os_waiting_tasks 25
 - Sys.dm_exec_requests 26
 - Querying sys.dm_exec_requests 28
 - Combining DMVs to Detect Waits Right Now 29
 - Viewing Wait Statistics Using Perfmon 33
 - Capturing Wait Statistics Using Extended Events 34
 - Capture Wait Statistics Information for a Specific Query 35
 - Summary 45
- **Chapter 3: Building a Solid Baseline 47**
 - What Are Baselines? 48
 - Visualizing Your Baselines 49
 - Baselines Types and Statistics 50
 - Baseline Pitfalls 51
 - Too Much Information 52
 - Know Your Metrics 52
 - Focus on the Big Measurement Changes 52
 - Use Fixed Intervals 52
 - Building a Baseline for Wait Statistics Analysis 52
 - Reset Capture Method 56
 - Delta Capture Method 56
 - Using SQL Server Agent to Schedule Measurements 58
 - Wait Statistics Baseline Analysis 60
 - Summary 67

■ **Part II: Wait Types..... 69**

■ **Chapter 4: CPU-Related Wait Types 71**

 CXPACKET..... 71

 What Is the CXPACKET Wait Type? 71

 Lowering CXPACKET Wait Time by Tuning the Parallelism Configuration 74

 Lowering CXPACKET Wait Time by Resolving Skewed Workloads 77

 CXPACKET Summary..... 78

 SOS_SCHEDULER_YIELD 78

 What Is the SOS_SCHEDULER_YIELD Wait Type? 79

 Lowering SOS_SCHEDULER_YIELD Waits 82

 SOS_SCHEDULER_YIELD Summary 85

 THREADPOOL..... 85

 What Is the THREADPOOL Wait Type? 86

 THREADPOOL Example 88

 Gaining Access to Our SQL Server during THREADPOOL Waits 91

 Lowering THREADPOOL Waits Caused by Parallelism 92

 Lowering THREADPOOL Waits Caused by User Connections..... 94

 THREADPOOL Summary..... 96

■ **Chapter 5: IO-Related Wait Types 97**

 ASYNC_IO_COMPLETION 97

 What Is the ASYNC_IO_COMPLETION Wait Type? 98

 ASYNC_IO_COMPLETION Example..... 98

 Lowering ASYNC_IO_COMPLETION Waits 99

 ASYNC_IO_COMPLETION Summary 103

 ASYNC_NETWORK_IO 103

 What Is the ASYNC_NETWORK_IO Wait Type? 103

 ASYNC_NETWORK_IO Example 104

- Lowering ASYNC_NETWORK_IO Waits 104
- ASYNC_NETWORK_IO Summary 106
- CMEMTHREAD 106**
 - What Is the CMEMTHREAD Wait Type? 106
 - Lowering CMEMTHREAD Waits 107
 - CMEMTHREAD Summary 108
- IO_COMPLETION 109**
 - What Is the IO_COMPLETION Wait Type? 109
 - IO_COMPLETION Example 109
 - Lowering IO_COMPLETION Waits 111
 - IO_COMPLETION Summary 111
- LOGBUFFER & WRITELOG 111**
 - What Are the LOGBUFFER & WRITELOG Wait Types? 111
 - LOGBUFFER & WRITELOG Example 113
 - Lowering LOGBUFFER & WRITELOG Waits 115
 - LOGBUFFER & WRITELOG Summary 116
- RESOURCE_SEMAPHORE 116**
 - What Is the RESOURCE_SEMAPHORE Wait Type? 116
 - RESOURCE_SEMAPHORE Example 117
 - Lowering RESOURCE_SEMAPHORE Waits 121
 - RESOURCE_SEMAPHORE Summary 122
- RESOURCE_SEMAPHORE_QUERY_COMPILE 122**
 - What Is the RESOURCE_SEMAPHORE_QUERY_COMPILE Wait Type? 122
 - RESOURCE_SEMAPHORE_QUERY_COMPILE Example 124
 - Lowering RESOURCE_SEMAPHORE_QUERY_COMPILE Waits 127
 - RESOURCE_SEMAPHORE_QUERY_COMPILE Summary 128

SLEEP_BPOOL_FLUSH..... 128

 What Is the SLEEP_BPOOL_FLUSH Wait Type? 128

 SLEEP_BPOOL_FLUSH Example 130

 Lowering SLEEP_BPOOL_FLUSH Waits..... 132

 SLEEP_BPOOL_FLUSH Summary..... 133

WRITE_COMPLETION 133

 What Is the WRITE_COMPLETION Wait Type? 133

 WRITE_COMPLETION Example..... 133

 Lowering WRITE_COMPLETION Waits 134

 WRITE_COMPLETION Summary..... 134

■ **Chapter 6: Backup-Related Wait Types 135**

 BACKUPBUFFER..... 136

 What Is the BACKUPBUFFER Wait Type? 136

 BACKUPBUFFER Example 138

 Lowering BACKUPBUFFER Waits 139

 BACKUPBUFFER Summary 140

 BACKUPIO 140

 What Is the BACKUPIO Wait Type? 140

 BACKUPIO Example..... 141

 Lowering BACKUPIO Waits..... 142

 BACKUPIO Summary..... 142

 BACKUPTHREAD 142

 What Is the BACKUPTHREAD Wait Type?..... 142

 BACKUPTHREAD Example 143

 Lowering BACKUPTHREAD Waits 144

 BACKUPTHREAD Summary 144

■ Chapter 7: Lock-Related Wait Types 145

- Introduction to Locking & Blocking 146
 - Lock Modes and Compatibility..... 146
 - Locking Hierarchy 148
 - Isolation Levels..... 149
 - Querying Lock Information 152
- LCK_M_S..... 155**
 - What Is the LCK_M_S Wait Type? 156
 - LCK_M_S Example 156
 - Lowering LCK_M_S Waits..... 157
 - LCK_M_S Summary..... 158
- LCK_M_U..... 158**
 - What Is the LCK_M_U Wait Type? 158
 - LCK_M_U Example 160
 - Lowering LCK_M_U Waits..... 161
 - LCK_M_U Summary..... 161
- LCK_M_X..... 162**
 - What Is the LCK_M_X Wait Type? 162
 - LCK_M_X Example 162
 - Lowering LCK_M_X Waits..... 163
 - LCK_M_X Summary..... 163
- LCK_M_I[xx] 163**
 - What Is the LCK_M_I[xx] Wait Type? 164
 - LCK_M_I[xx] Example..... 164
 - Lowering LCK_M_I[xx] Waits 166
 - LCK_M_I[xx] Summary..... 167

LCK_M_SCH_S & LCK_M_SCH_M.....	167
What Are the LCK_M_SCH_S & LCK_M_SCH_M Wait Types?.....	167
LCK_M_SCH_S & LCK_M_SCH_M Example	168
Lowering LCK_M_SCH_S & LCK_M_SCH_M Waits	170
LCK_M_SCH_S & LCK_M_SCH_M Summary	170
■ Chapter 8: Latch-Related Wait Types.....	171
Introduction to Latches	171
Latch Modes	172
Latch Waits	173
Sys.dm_os_latch_stats	175
Page-Latch Contention	176
PAGELATCH_[xx].....	180
What Is the PAGELATCH_[xx] Wait Type?	180
PAGELATCH_[xx] Example	180
Lowering PAGELATCH_[xx] Waits.....	184
PAGELATCH_[xx] Summary.....	189
LATCH_[xx]	189
What Is the LATCH_[xx] Wait Type?	190
LATCH_[xx] Example.....	191
Lowering LATCH_[xx] Waits	196
LATCH_[xx] Summary	196
PAGEIOLATCH_[xx].....	197
What Is the PAGEIOLATCH_[xx] Wait Type?	197
PAGEIOLATCH_[xx] Example	199
Lowering PAGEIOLATCH_[xx] Waits	200
PAGEIOLATCH_[xx] Summary	205

■ **Chapter 9: High-Availability and Disaster-Recovery Wait Types ... 207**

DBMIRROR_SEND..... 208

 What Is the DBMIRROR_SEND Wait Type? 211

 DBMIRROR_SEND Example 211

 Lowering DBMIRROR_SEND Waits..... 213

 DBMIRROR_SEND Summary..... 214

HADR_LOGCAPTURE_WAIT & HADR_WORK_QUEUE..... 214

 What Are the HADR_LOGCAPTURE_WAIT & HADR_WORK_QUEUE Wait Types?... 214

 HADR_LOGCAPTURE_WAIT & HADR_WORK_QUEUE Summary..... 216

HADR_SYNC_COMMIT 217

 What Is the HADR_SYNC_COMMIT Wait Type? 217

 HADR_SYNC_COMMIT Example..... 218

 Lowering HADR_SYNC_COMMIT Waits..... 220

 HADR_SYNC_COMMIT Summary..... 223

REDO_THREAD_PENDING_WORK..... 223

 What Is the REDO_THREAD_PENDING_WORK Wait Type? 223

 REDO_THREAD_PENDING_WORK Summary..... 225

■ **Chapter 10: Preemptive Wait Types..... 227**

PREEMPTIVE_OS_ENCRYPTMESSAGE & PREEMPTIVE_OS_DECRYPTMESSAGE..... 228

 What Are the PREEMPTIVE_OS_ENCRYPTMESSAGE & PREEMPTIVE_OS_DECRYPTMESSAGE Wait Types? 228

 PREEMPTIVE_OS_ENCRYPTMESSAGE & PREEMPTIVE_OS_DECRYPTMESSAGE Example..... 229

 Lowering PREEMPTIVE_OS_ENCRYPTMESSAGE & PREEMPTIVE_OS_DECRYPTMESSAGE Waits 235

 PREEMPTIVE_OS_ENCRYPTMESSAGE & PREEMPTIVE_OS_DECRYPTMESSAGE Summary 235

PREEMPTIVE_OS_WRITEFILEGATHER.....	236
What Is the PREEMPTIVE_OS_WRITEFILEGATHER Wait Type?	236
PREEMPTIVE_OS_WRITEFILEGATHER Example	236
Lowering PREEMPTIVE_OS_WRITEFILEGATHER Waits.....	237
PREEMPTIVE_OS_WRITEFILEGATHER Summary.....	238
PREEMPTIVE_OS_AUTHENTICATIONOPS	239
What Is the PREEMPTIVE_OS_AUTHENTICATIONOPS Wait Type?.....	239
PREEMPTIVE_OS_AUTHENTICATIONOPS Example.....	240
Lowering PREEMPTIVE_OS_AUTHENTICATIONOPS Waits	241
PREEMPTIVE_OS_AUTHENTICATIONOPS Summary	242
PREEMPTIVE_OS_GETPROCADDRESS	243
What Is the PREEMPTIVE_OS_GETPROCADDRESS Wait Type?	243
PREEMPTIVE_OS_GETPROCADDRESS Example.....	245
Lowering PREEMPTIVE_OS_GETPROCADDRESS Waits.....	245
PREEMPTIVE_OS_GETPROCADDRESS Summary.....	246
■ Chapter 11: Background and Miscellaneous Wait Types.....	247
CHECKPOINT_QUEUE	248
What Is the CHECKPOINT_QUEUE Wait Type?	248
CHECKPOINT_QUEUE Summary.....	251
DIRTY_PAGE_POLL	251
What Is the DIRTY_PAGE_POLL Wait Type?.....	251
DIRTY_PAGE_POLL Summary	253
LAZYWRITER_SLEEP.....	253
What Is the LAZYWRITER_SLEEP Wait Type?	253
LAZYWRITER_SLEEP Summary.....	255
MSQL_XP.....	255
What Is the MSQL_XP Wait Type?	255
MSQL_XP Example	256

Lowering MSQL_XP Waits.....	257
MSQL_XP Summary.....	257
OLEDB.....	257
What Is the OLEDB Wait Type?.....	257
OLEDB Example	258
Lowering OLEDB Waits	258
OLEDB Summary	259
TRACEWRITE.....	259
What Is the TRACEWRITE Wait Type?.....	260
TRACEWRITE Example	260
Lowering TRACEWRITE Waits.....	263
TRACEWRITE Summary	266
WAITFOR.....	266
What Is the WAITFOR Wait Type?	266
WAITFOR Example	267
WAITFOR Summary.....	267
■ Chapter 12: In-Memory OLTP–Related Wait Types.....	269
Introduction to In-Memory OLTP.....	270
CFPs.....	270
Isolation	273
Transaction Log Changes.....	274
WAIT_XTP_HOST_WAIT	274
What Is the WAIT_XTP_HOST_WAIT Wait Type?.....	274
WAIT_XTP_HOST_WAIT Summary	278
WAIT_XTP_CKPT_CLOSE	279
What Is the WAIT_XTP_CKPT_CLOSE Wait Type?.....	279
WAIT_XTP_CKPT_CLOSE Summary	281

WAIT_XTP_OFFLINE_CKPT_NEW_LOG 281

 What Is the WAIT_XTP_OFFLINE_CKPT_NEW_LOG Wait Type?..... 281

 WAIT_XTP_OFFLINE_CKPT_NEW_LOG Summary 283

■ **Appendix I: Example SQL Server Machine Configurations 285**

 Default Test Machine 285

 HA/DR Test Machines 286

 In-Memory OLTP Test Machine 287

■ **Appendix II: Spinlocks..... 289**

■ **Appendix III: Latch Classes 293**

Index..... 301

About the Author



Enrico van de Laar is a database technology specialist with over a decade of experience with SQL Server. Enrico works for a Dutch IT company that specializes in building IT solutions for the Dutch healthcare sector, where he is responsible for maintaining and optimizing the databases of some of the largest hospitals in the Netherlands. Enrico is also the founder of DotNine SQL Server Training, a company that specializes in Microsoft SQL Server-related training.

Enrico is a SQL Server MVP, and his main specialties within SQL Server are performance tuning and database architecture.

You can contact Enrico through his Twitter handle @evdlaar, or by sending him an email at enrico@dotnine.nl.

About the Technical Reviewer



Chris Presley is a SQL programmer, database administrator (DBA), and consultant who has been working with data in IT since 1998. Chris shared his expertise and passion with budding DBAs as a former instructor at Conestoga College in Kitchener, Ontario. As the Director of Consulting, Chris currently leads a team of over 20 data consultants at Pythian, a global enterprise IT consulting and managed services company. Chris's particular areas of expertise are in disaster recovery, database administration, architecture, troubleshooting multi-system issues, and integration services. He's currently expanding his expertise into Open Source databases, data science, and big data.

He speaks at user groups and other industry events. You can follow him on Twitter at @DBAChris and find him on LinkedIn: <https://www.linkedin.com/in/chrispresley>. His blog can be found at <http://www.pythian.com/blog/author/presley/>.

Acknowledgments

First of all, I would like to thank my wife, Margreet, and my children, Sander, Thijs, and Emma. Writing a book requires an immense investment in time and dedication. Without the understanding and support of my family I would never have been able to write this book.

Second of all, I want to thank the great people working at Apress for believing in me and giving me the chance to write this book. Jonathan and Jill, thanks for your patience and help with all the questions I had!

A big thanks goes to Chris Presley for being an amazing technical editor. Your great feedback and knowledge about SQL Server helped me a lot during the writing of this book.

I am lucky to be working for one of the best IT companies in the Netherlands—PinkRocade Healthcare—that helps and supports me with every step I take. To all my colleagues, thanks for your patience while I endlessly talked about either writing this book or SQL Server in general!

Last, but certainly not least, I would like to thank the complete SQL Server family. It is a great honor to be part of such a friendly and helpful group, you never stop inspiring and motivating me.

Introduction

Performance is a hot issue on a lot of database implementations. Many businesses run into performance-related issues when their databases experience more load or grow larger in size. There are many methods available for increasing the performance of your SQL Server(s) on all types of levels. Many of these performance-optimization methods we consider best-practice, like running index maintenance to make sure fragmented indexes don't slow down your queries, or updating statistics so the SQL Server Database Engine has the correct information to generate a good execution plan. Besides these database maintenance methods, you can also choose to dive a little bit deeper into specific query performance troubleshooting, optimizing queries by making sure expensive operators are replaced by less expensive ones, for instance. And of course there is always the "sledgehammer" approach, replacing your current hardware for newer, better performing hardware, hoping that will solve the performance issues you are experiencing.

No matter what approach you choose to optimize or troubleshoot SQL Server performance, there are always two common resources involved: time and money. Ideally we want to spend as little time and money as possible while we are working on increasing performance. Knowing where to focus your time and money is very important. If you can find the source of the performance problem and resolve it at that level, you can save a lot of time and money that you would have spent on analyzing symptoms.

In a way, we can compare our search for the heart of our performance issues with a medical examination. Instead of giving out different types of medication until something actually works, a physician is always trying to find the source of the problem so he or she can prescribe the right medication that works best for that specific condition without causing side effects. The same approach works for SQL Server. Implementing all types of possible solutions without looking at the real source of the problem will probably not solve the real underlying issue (unless you're really lucky) and can possibly make matters worse.

This is where Wait Statistics can help. Wait Statistics are generated and maintained at the heart of the SQL Server Database Engine where queries are being executed, giving valuable insight into what is slowing down your queries. There are 771 different types of Wait Statistics in the latest edition of SQL Server (SQL Server 2014), and with every edition that number grows as new features are introduced or existing features are modified or expanded. That is a lot of information that is freely available to help you troubleshoot!

This book is my attempt to help you understand SQL Server Wait Statistics. It will go into detail how Wait Statistics are being generated and how you can use that information to optimize, or troubleshoot, the performance of your SQL Server installation. I will also describe specific Wait Statistics and give you pointers on how you can resolve problems yourself. In the case of this book, I personally believe the journey is more important than reaching the destination. For that reason, I spend more time describing and explaining

what is causing the specific Wait Types to occur than I do writing down every possible way you can lower their Wait Times. If you understand why a Wait Type is generated, and to what part of SQL Server it is related, you will quickly discover methods of your own to lower their Wait Times.

Because of the sheer number of different Wait Statistics it is sadly impossible to describe and discuss all of them. For this reason, I had to make a selection of Wait Statistics to include in this book. The way I did this was by gathering Wait Statistics information off many different SQL Server installations and selecting the most common or most performance-degrading ones, resulting in a selection of 45 different (or grouped) Wait Types.

Why This Book?

As I wrote in the introduction, the goal of this book is to give you a deeper understanding of SQL Server Wait Statistics and also to describe various Wait Statistics in detail. For this reason this book has been split up into two parts, Part I describing the foundation of Wait Statistics analysis and Part II describing various specific Wait Statistics in detail. I tried to categorize the Wait Statistics in Part II by the part of the system they affect (i.e., CPU, Memory, etc.). Some Wait Statistics aren't that easily categorized, since they can affect multiple system parts. In those cases, I tried to categorize them with the part they have the most effect on.

Part I: Foundations of Wait Statistics Analysis

- Chapter 1: "Wait Statistics Internals" starts off with a brief history of SQL Server Wait Statistics and a look at the SQLOS architecture. Because Wait Statistics have a close relationship with the processor(s) of your system, we will discuss schedulers, tasks, and worker threads in detail.
- Chapter 2: "Querying SQL Server Wait Statistics" introduces the various way you can access the Wait Statistics information inside the SQL Server Database Engine by using DMVs, Extended Events, and Perfmon.
- Chapter 3: "Building a Solid Baseline" covers the importance of building and using a baseline for performance troubleshooting. Baselines are especially important when analyzing Wait Statistics, and I will show you examples of how you can create one for Wait Statistics analysis.

Part II: Wait Types

- Chapter 4: "CPU-Related Wait Types" introduces Wait Types that are CPU related.
- Chapter 5: "IO-Related Wait Types" describes Wait Types that are IO related.
- Chapter 6: "Backup-Related Wait Types" presents the Wait Types that are related to backup events.

- Chapter 7: “Lock-Related Wait Types” starts off with a short introduction to locking and blocking in SQL Server before diving into the lock-related Wait Types.
- Chapter 8: “Latch-Related Wait Types” starts off with a close look at latches, describing what they are and how they work inside SQL Server. After this introduction we are ready to take a look at different latch-related Wait Types.
- Chapter 9: “High-Availability and Disaster-Recovery Wait Types” describes Wait Types that are related to the various HA & DR configurations available in SQL Server.
- Chapter 10: “Preemptive Wait Types” presents different Wait Types that have a direct relationship with the operating system of your SQL Server.
- Chapter 11: “Background and Miscellaneous Wait Types” includes various Wait Types that are being generated by the SQLOS as a background process. This is also the chapter where we will describe Wait Types that couldn’t be categorized in one of the above categories.
- Chapter 12: “In-Memory OLTP-Related Wait Types” describes some of the new Wait Types that were introduced in SQL Server 2014 with the release of the In-Memory OLTP feature.
- Appendix I: “Example SQL Server Machine Configuration” describes the configuration of the Virtual Machines I used during the book.
- Appendix II: “Spinlocks” explains the working of so-called lightweight synchronization primitives called spinlocks and the impact they can have on your SQL Server configuration.
- Appendix III: “Latch Classes” contains a list of a large portion of different latch classes inside SQL Server. The list is a combination of information from Books Online and additional information about the specific latch class.

Word of Warning

The way we access Wait Statistics information in this book is mostly by SQL queries. For this reason this book has quite a lot of lines of SQL code. Most of the queries in Part I of this book deal with gathering or capturing Wait Statistics using Dynamic Management Views (or DMVs) and as such are, unless stated otherwise, harmless. Some of the queries in Part II, however, are written to actually harm performance so specific Wait Types can be demonstrated. Please keep this in mind when you plan to use some of the scripts and test them thoroughly on a test system.

PART I



Foundations of Wait Statistics Analysis

CHAPTER 1



Wait Statistics Internals

SQL Server Wait Statistics are an important tool you can use to analyze performance-related problems or to optimize your SQL Server's performance. They are, however, not that well known to many database administrators or developers. I believe this has to do with their relatively complex nature, the sheer volume of the different types of Wait Statistics, and the lack of documentation for many types of Wait Statistics. Wait Statistics are also directly related to the SQL Server you are analyzing them on, which means that it is impossible to compare the Wait Statistics of Server A to the Wait Statistics of Server B, even if they had an identical hardware and database configuration. Every configuration option, from the hardware firmware level to the configuration of the SQL Server Native Client on the client computers, will have an impact on the Wait Statistics!

For the reasons just mentioned, I firmly believe we should start with the foundation and internals of SQL Server Wait Statistics so you can get familiar with how they are generated, how you can access them, and how you can use them for performance troubleshooting. This approach will get you ready for Part II of this book, where we will examine specific Wait Statistics.

In this chapter we will take a brief look at the history of Wait Statistics through the various versions of SQL Server. Following that, we will take a close look at the SQL Operating System, or SQLOS. The architecture of the SQLOS is closely tied to Wait Statistics and to performance troubleshooting in general. The rest of the chapter is dedicated to one of the most important aspects of Wait Statistics—thread scheduling.

Before we begin with the foundation and internals of SQL Server Wait Statistics, I would like to mention a few things related to the terminology used when discussing Wait Statistics. In the introduction of this book and the paragraphs above, I only mentioned the term Wait Statistics. The sentence “compare the Wait Statistics of Server A to the Wait Statistics of Server B” is actually wrong, since we can only compare the *Wait Time* (the total time we have been waiting on a resource) of a specific *Wait Type* (the specific Wait Type related to the resource we are waiting on). From this point on, when I use the term *Wait Statistics* I mean the concept of Wait Statistics, and I will use the correct terms Wait Time and Wait Type where appropriate.

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-1-4842-1139-7_1](https://doi.org/10.1007/978-1-4842-1139-7_1)) contains supplementary material, which is available to authorized users.

A Brief History of Wait Statistics

SQL Server has been around for quite some time now; the first release of SQL Server dates back to 1989 and was released for the OS/2 platform. Until SQL Server 6.0, released in 1995, Microsoft worked together with Sybase to develop SQL Server. In 1995, however, Microsoft and Sybase went their own ways. Microsoft and Sybase stayed active in the database world (SAP actually acquired Sybase in 2010), and in 2014 Microsoft released SQL Server 2014 and SAP released SAP Sybase ASE 16, both relational enterprise-level database systems.

Between SQL Server 6.0 and SQL Server 2014, so many things have changed that you simply cannot compare the two versions. One thing that hasn't changed in all these years is Wait Statistics. In one way or another, SQL Server stores information about its internal processes, and even though the way we access that information has changed over the years, Wait Statistics remain an important part of the internal logging process.

In early versions of SQL Server we needed to access the Wait Statistics using undocumented commands. Figure 1-1 shows how you would query Wait Statistics information in SQL Server 6.5 using the DBCC command.

Wait Statistic	Value
DataserVSea waits/tran	0.0
Rlock requests/tran	21.5882
Rlock waits/tran	0.0
Rlock millisecond/tran	0.0
PageLock requests/tran	15.1765
PageLock waits/tran	0.0
PageLock millisecond/tran	0.0
TableLock requests/tran	2.0
TableLock waits/tran	0.0
TableLock millisecond/tran	0.0
ExtLock requests/tran	0.588235
ExtLock waits/tran	0.0
ExtLock millisecond/tran	0.0
Upwails/tran	0.382353
AppndLog waits/tran	0.0
Upwait millisecond/tran	0.0
WriteLog waits/tran	0.0
WriteLog millisecond/tran	0.0
Budget diskreads/tran	0.0
Budget millisecond/tran	0.0
Total millisecond/tran	0.0

(21 row(s) affected)

DBCC execution completed. If DBCC printed error messages, see your System Administrator.

Figure 1-1. SQL Server Wait Statistics in SQL Server 6.5

One of the big changes that was introduced in SQL Server 2005 was the conversion of many internal functions and commands into Dynamic Management Views (DMV), including Wait Statistics information. This made it far easier to query and analyze the information returned by functions or commands. A new way of performance analysis was born with the release of the SQL Server 2005 Microsoft whitepaper “SQL Server 2005 Waits and Queues” by Tom Davidson.

In the various releases of SQL Server the amount of different Wait Types grew exponentially whenever new features or configuration options were introduced. If you take a good look at Figure 1-1 you will notice that 21 different Wait Types were returned. Figure 1-2 shows the amount of Wait Types, as the number of rows returned, available in SQL Server 2014.

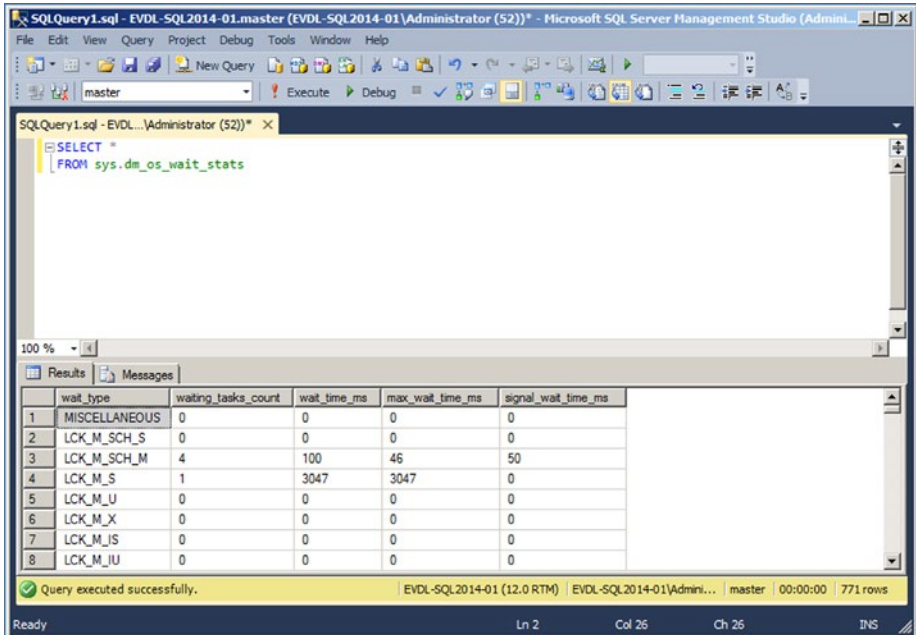


Figure 1-2. SQL Server Wait Statistics in SQL Server 2014

Those 771 rows are all different Wait Types and hold wait information for different parts of the SQL Server engine. The number of Wait Types will continue to grow in future SQL Server releases, as new features are introduced or existing features are changed. Thankfully there is a lot more information available about Wait Statistics now than there was in SQL Server 6.5!

The SQLLOS

The world of computer hardware changes constantly. Every year, or in some cases every month, we manage to put more cores inside a processor, increase the memory capacity of mainboards, or introduce entirely new hardware concepts like PCI-based persistent flash storage. Database Management Systems (or DBMSs) are always one of the first types of applications that want to take advantage of new hardware trends. Because of the fast-changing nature of hardware and the need to utilize new hardware options as soon as they become available, the SQL Server team decided to change the SQL Server platform layer in SQL Server 2005.

Before SQL Server 2005, the platform layer of SQL Server was pretty restricted, and many operations were performed by the operating system. This meant that it was difficult for SQL Server to keep up with the fast-changing world of server hardware, as changing a complete operating system in order to utilize faster hardware or new hardware features is a time-consuming and complex operation.

Figure 1-3 shows the (simplified) architecture of SQL Server before the introduction of the SQLOS in SQL Server 2005.

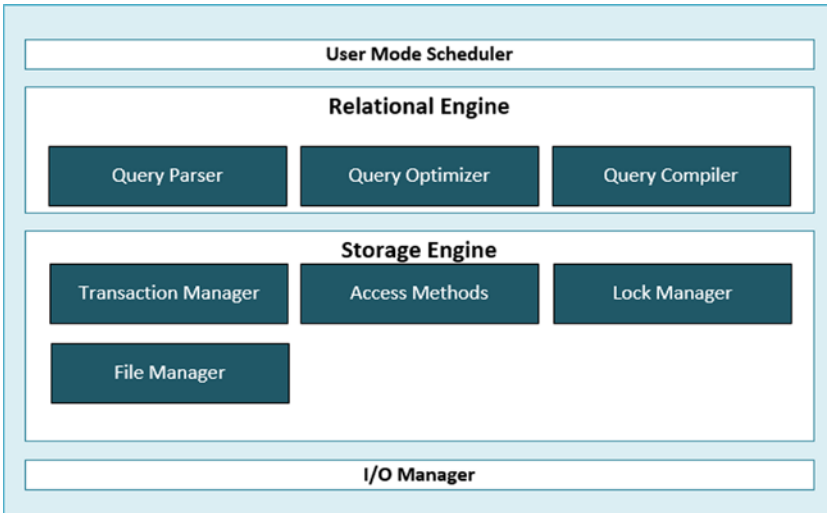


Figure 1-3. SQL Server architecture before the introduction of the SQLOS

SQL Server 2005 introduced one of the biggest changes to the SQL Server engine seen to this day, the SQLOS. This is a completely new platform layer that functions as a user-level operating system. This new operating system has made it possible to fully utilize current and future hardware and has enabled features like advanced parallelism. The SQLOS is highly configurable and adjusts itself to the hardware it is running on, thus making it perfectly scalable for high-end or low-end systems alike.

Figure 1-4 shows the (simplified) architecture of SQL Server 2005, including the SQLOS layer.

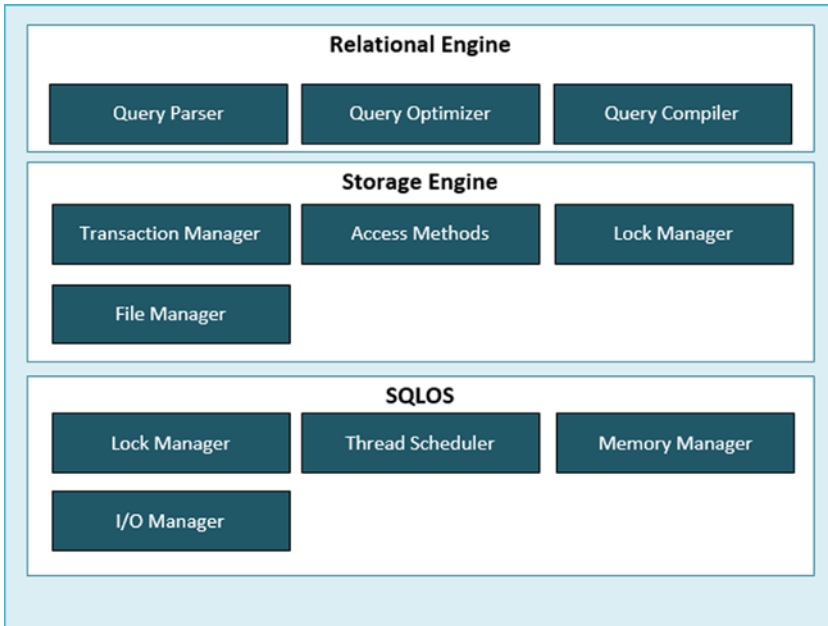


Figure 1-4. SQL Server 2005 architecture

The SQLOS changed the way SQL Server accesses processor resources by introducing schedulers, tasks, and worker threads. This gives the SQLOS greater control of how work should be completed by the processors. The Windows operating system uses a *preemptive scheduling* approach. This means that Windows will give every process that needs processor time a priority and fixed slice of time, or a *quantum*. This process priority is calculated from a number of variables like resource usage, expected runtime, current activity, and so forth. By using preemptive scheduling, the Windows operating system can choose to interrupt a process when a process with a higher priority needs processor time. This way of scheduling can have a negative impact on processes generated by SQL Server, since those processes could easily be interrupted by higher priority ones, including those of other applications. For this reason, the SQLOS uses its own (cooperative) non-preemptive scheduling mechanism, making sure that Windows processes cannot interrupt SQLOS processes.

SQL Server 7 and SQL Server 2000 also used non-preemptive scheduling using User Mode Scheduling (UMS). SQLOS brought many more system components closer together, thus enabling better performance and scalability.

There are some exceptions when the SQLOS cannot use non-preemptive scheduling, for instance, when the SQLOS needs to access a resource through the Windows operating system. We will discuss these exceptions later in this book in the “Preemptive Wait Types” chapter.

Schedulers, Tasks, and Worker Threads

Because the SQLOS uses a different method to execute requests than the Windows operating system uses, SQL Server introduced a different way to schedule processor time using schedulers, tasks, and worker threads. Figure 1-5 shows the different parts of SQL Server scheduling and how they relate to each other.




Figure 1-5. SQL Server scheduling

Sessions

A *session* is the connection a client has to the SQL Server it is connected to (after it has been successfully authenticated). We can easily access session information by querying the **sys.dm_exec_sessions** DMV using the following query:

```
SELECT * FROM sys.dm_exec_sessions;
```

Generally speaking, user sessions will have a `session_id` higher than 50; everything lower is reserved for internal SQL Server processes. However, on very busy servers there is a possibility that SQL Server needs to use a `session_id` higher than 50. If you are only interested in information about user-initiated sessions, it is better to filter the results of the `sys.dm_exec_sessions` DMV using the `is_user_process` column instead of filtering on a `session_id` greater than 50. The following query will only return user sessions and will filter out the internal system sessions:

```
SELECT * FROM sys.dm_exec_sessions
WHERE is_user_process = 1;
```

Figure 1-6 shows a small part of the results of this query.

	session_id	login_time	host_name	program_name	host_process_id	client_version
1	51	2015-02-10 11:12:43.123	EVDL-SQL2014-01	Microsoft SQL Server Management Studio	924	7
2	52	2015-02-10 11:12:52.807	EVDL-SQL2014-01	Microsoft SQL Server Management Studio - Query	924	7

Figure 1-6. *sys.dm_exec_sessions* results

There are many more columns returned by the **sys.dm_exec_sessions** DMV that will give us information about the specific session. Some of the more interesting columns that deserve some extra explanation are the `host_process_id`, which is the Process ID (or PID) of the client program connected to the SQL Server. The `cpu_time` column will give you information about the processor time (in milliseconds) the session has used since it was first established. The `memory_usage` column displays the amount of memory used by the session. This is not the amount in MB or KB, but the number of 8 KB pages used. Another column I would like to highlight is the `status` column. This will show you if the session has any active requests. The most common values of the `status` column are “Running,” which indicates that one or more requests are currently being processed from this session, and “Sleeping,” which means no requests are currently being processed from this session.

Requests

A *request* is the SQL Server execution engine’s representation of a query submitted by a session. Again, we can use a DMV to query information about a request; in this case, we can run a query against the **sys.dm_exec_requests** DMV like the query below:

```
SELECT * FROM sys.dm_exec_requests;
```

Figure 1-7 shows a portion of the results of this query.

	session_id	request_id	start_time	status	command	sj_handle	statement_start_offset	statement_end_offset	plan_handle
32	32	0	2015-02-10 12:40:54.980	sleeping	TASK MANAGER	NULL	NULL	NULL	NULL
33	33	0	2015-02-10 12:40:54.980	sleeping	TASK MANAGER	NULL	NULL	NULL	NULL
34	34	0	2015-02-10 12:37:49.457	sleeping	TASK MANAGER	NULL	NULL	NULL	NULL
35	35	0	2015-02-10 12:37:49.457	sleeping	TASK MANAGER	NULL	NULL	NULL	NULL
36	36	0	2015-02-10 12:40:54.980	sleeping	TASK MANAGER	NULL	NULL	NULL	NULL
37	39	0	2015-02-10 12:37:49.457	sleeping	TASK MANAGER	NULL	NULL	NULL	NULL
38	53	0	2015-02-10 12:43:58.690	running	SELECT	0x02000...	0	66	0x060001...

Figure 1-7. *sys.dm_exec_requests* results

The **sys.dm_exec_requests** DMV is an incredibly powerful tool to use when you are troubleshooting any performance-related issues. The reason for this is that it has a lot of information about the actual queries being executed and can help you detect performance bottlenecks relatively quickly. Because the **sys.dm_exec_requests** DMV also displays Wait Statistics–related information, we will take a thorough look at it in Chapter 2, “Querying Wait Statistics.”

Tasks

Tasks represent the actual work that needs to be performed by the SQLOS, but they do not perform any work themselves. When a request is received by SQL Server, one or more tasks will be created to fulfill the request. The amount of tasks that get generated for a request depends on if the query request is being performed using parallelism or if it’s being run serially.

We can use the **sys.dm_os_tasks** DMV to query the task information, like I did in the example query below:

```
SELECT * FROM sys.dm_os_tasks;
```

Figure 1-8 shows a part of the results of the query.

	task_address	task_state	context_switches_count	pending_io_count	pending_io_byte_count	pending_io_byte_average
1	0x000000017F0164E8	SUSPENDED	764	0	0	0
2	0x000000017F0168C8	RUNNING	4	0	0	0
3	0x000000017F016CA8	SUSPENDED	2260	0	0	0
4	0x000000017F017088	SUSPENDED	8941871	0	0	0
5	0x000000017F017848	SUSPENDED	2	0	0	0
6	0x000000017F017468	SUSPENDED	142529	0	0	0
7	0x000000017F017C28	SUSPENDED	537	43	0	0

Figure 1-8. *sys.dm_os_tasks* results

When you query the **sys.dm_os_tasks** DMV you will discover it will return many results, even on servers that have no user activity. This is because SQL Server uses tasks for its own processes as well; you can identify those by looking at the `session_id` column.

There are some interesting columns in this DMV that are worth exploring to see the relations between the different DMVs. The `task_address` column will show you the memory address of the task. The `session_id` will return the ID of the session that has requested the task, and the `worker_address` will hold the memory address of the worker thread associated with the task.

Worker Threads

Worker threads are where the actual work for the request is being performed. Every task that gets created will get a worker thread assigned to it, and the worker thread will then perform the actions requested by the task.

A worker thread will actually not perform the work itself; it will request a thread from the Windows operating system to perform the work for it. For the sake of simplicity, and the fact the actual Windows thread runs outside the SQLOS, I have left this step out of Figure 1-5. You can access information about the Windows operating system threads by querying **sys.dm_os_threads** if you are interested.

When a task requests a worker thread SQL Server will look for an idle worker thread and assign it to the task. In the case when no idle worker thread can be located and the maximum amount of worker threads has been reached, the request will be queued until a worker thread finishes its current work and becomes available.

There is a limit to the number of worker threads SQL Server has available for processing requests. This number will be automatically calculated and configured by SQL Server during the installation. We can also calculate the maximum amount of worker threads ourselves using these formulas:

- 32-bit system with less than, or equal to, 4 logical processors:
256 worker threads
- 32-bit system with more than 4 logical processors:
 $256 + ((\text{number of logical processors} - 4) * 8)$
- 64-bit system with less than, or equal to, 4 logical processors:
512 worker threads
- 64-bit system with more than 4 logical processors:
 $512 + ((\text{number of logical processors} - 4) * 16)$
- Example: If we have a 64-bit system with 16 processors (or cores) we can calculate the maximum number of worker threads using the formula, $512 + ((16 - 4) * 16)$, which would give us a maximum of 704 worker threads.

The amount of worker threads can be changed from the default of **0** (which means SQL Server sets the amount of max worker threads using the formulas above when it starts) by changing the **max worker threads** options in your SQL Server’s properties, as illustrated by Figure 1-9.

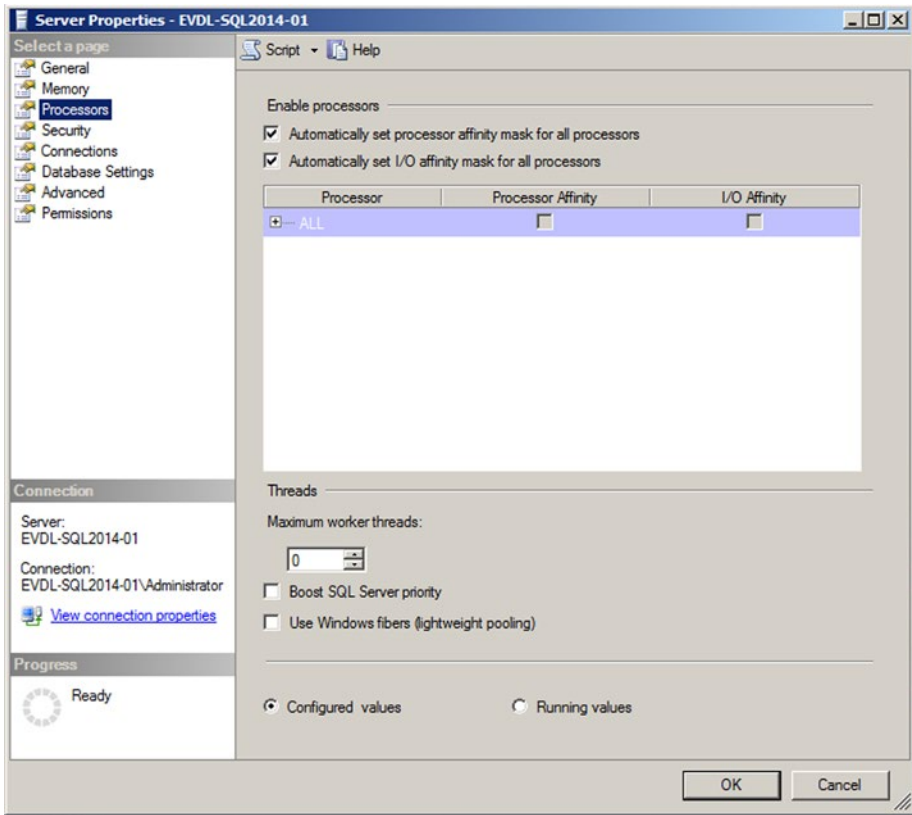


Figure 1-9. Processors page in the Server Properties

Generally speaking, there should be no need to change the **max worker threads** option, and my advice is to leave the setting alone, as it should only be changed in very specific cases (I will discuss one of those potential cases in Chapter 4 when we talk about THREADPOOL waits).

One thing to keep in mind is that worker threads require memory to work. For 32-bit systems this is 512 KB for every worker thread; 64-bit systems will need 2048 KB for every worker thread. Thus, changing the number of worker threads can potentially impact the memory requirements of SQL Server. This does not mean you need a massive amount of memory just for your worker threads—SQL Server will automatically destroy worker threads if they have been idle for 15 minutes or if your SQL Server is under heavy memory pressure.

SQL Server supplies us with a DMV to query information about the worker threads: **sys.dm_os_workers**. Figure 1-10 shows some of the results of this query:

```
SELECT * FROM sys.dm_os_workers;
```

	worker_address	status	is_preemptive	is_fiber	is_sick	is_in_cc_exception	is_fatal_exception	is_inside_catch
1	0x000000017F01A160	2	0	0	0	0	0	0
2	0x000000017F01C160	4	1	0	0	0	0	0
3	0x000000017F024160	2	0	0	0	0	0	0
4	0x000000017F200160	4	1	0	0	0	0	0
5	0x000000017F202160	4	1	0	0	0	0	0
6	0x000000017F21C160	4	1	0	0	0	0	0
7	0x000000017F250160	0	0	0	0	0	0	0
8	0x000000017F05A160	0	0	0	0	0	0	0

Figure 1-10. Results of querying **sys.dm_os_workers**

The **sys.dm_os_workers** DMV is a very large and complex DMV where many columns are marked as “Internal use only” by Microsoft. In this DMV the columns **task_address** and **scheduler_address** are available to link together the different DMVs we have discussed.

Worker threads go through different phases while they are being exposed to the processor, which we can view when we look at the state column in the **sys.dm_os_workers** DMV:

- **INIT.** The worker thread is being initialized by the SQLOS.
- **RUNNING.** The worker thread is currently performing work on a processor.
- **RUNNABLE.** The worker thread is ready to run on a processor.
- **SUSPENDED.** The worker thread is waiting for a resource.

The states the worker threads go through while performing their work are one of the main topics of this book. Every time a worker thread is not in the “RUNNING” state, it has to wait, and the SQLOS records this information into Wait Statistics, giving us valuable insight into what the worker thread has been waiting on and how long it has been waiting.

Schedulers

The *scheduler* component’s main task is to—surprise—schedule work, in the form of tasks, on the physical processor(s). When a task requests processor time it is the scheduler that assigns worker threads to that task so the request can get processed. It is also responsible for making sure worker threads cooperate with each other and yield the processor when their slice of time, or quantum, has expired. We call this *cooperative scheduling*. The need for worker threads to yield when their processor time has expired comes from the fact that a scheduler will only let one worker thread run on a processor at a time. If the worker threads didn’t need to yield, a worker thread could stay on the processor for an infinite amount of time, blocking all usage of that processor.

There is a one-on-one relation between processors and schedulers. If your system has two processors, each with four cores, there will be eight schedulers that the SQLOS can use to process user requests, each of them mapped to one of the logical processors.

We can access information about the schedulers by running a query against the **sys.dm_os_schedulers** DMV:

```
SELECT * FROM sys.dm_os_schedulers;
```

The results of the query are shown in Figure 1-11.

	scheduler_address	parent_node_id	scheduler_id	cpu_id	status	is_online	is_idle
1	0x000000017F120040	0	0	0	VISIBLE ONLINE	1	1
2	0x000000017F130040	0	1	1	VISIBLE ONLINE	1	0
3	0x000000017F140040	0	1048578	0	HIDDEN ONLINE	1	0
4	0x000000017F320040	64	1048576	0	VISIBLE ONLINE (DAC)	1	1
5	0x000000017F370040	0	1048579	1	HIDDEN ONLINE	1	1
6	0x000000017B980040	0	1048580	0	HIDDEN ONLINE	1	1
7	0x000000017B9E0040	0	1048581	1	HIDDEN ONLINE	1	1

Figure 1-11. *sys.dm_os_schedulers* query results

The SQL Server on which I ran this query has one processor with two cores, which means there should be two schedulers that can process my user requests. If we look at Figure 1-11, however, we notice there are more than two schedulers returned by the query. SQL Server uses its own schedulers to perform internal tasks, and those schedulers are also returned by the DMV and are marked “HIDDEN ONLINE” in the status column of the DMV. The schedulers that are available for user requests are marked as “VISIBLE ONLINE” in the DMV. There is also a special type of scheduler with the status “VISIBLE ONLINE (DAC).” This is a scheduler dedicated for use with the Dedicated Administrator Connection. This scheduler makes it possible to connect to SQL Server in situations where it is unresponsive; for instance, when there are no free worker threads available on the schedulers that process user requests.

We can view the number of worker threads a scheduler has associated with it by looking at the *current_workers_count* column. This number also includes worker threads that aren’t performing any work. The *active_workers_count* shows us the worker threads that are active on the specific scheduler. This doesn’t mean they are actually

running on the processor, as worker threads with states of “RUNNING,” “RUNNABLE,” and “SUSPENDED” also count toward this number. The `work_queue_count` is also an interesting column since it will give you insight into how many tasks are waiting for a free worker thread. If you see high numbers in this column, it might mean that you are experiencing CPU pressure.

Putting It All Together

All the parts of the SQL Server scheduling we have discussed so far are connected to each other, and every request passes through these same components. The following text is an example of how a query request would get processed.

A user connects to the SQL Server through an application. The SQL Server will create a session for that user after the login process is completed successfully. When the user sends a query to the SQL Server, a task and a request will be created to represent the unit of work that needs to be done. The scheduler will assign worker threads to the task so it can be completed.

To see all this information in SQL Server, we can join some of the DMVs we used. The query in Listing 1-1 will show you an example of how we can combine the different DMVs to get scheduling information about a specific session (in this case a session with an ID of 53).

Listing 1-1. Join the different DMVs together to query scheduling information

```
SELECT
    r.session_id AS 'Session ID',
    r.command AS 'Type of Request',
    qt.[text] AS 'Query Text',
    t.task_address AS 'Task Address',
    t.task_state AS 'Task State',
    w.worker_address AS 'Worker Address',
    w.[state] AS 'Worker State',
    s.scheduler_address AS 'Scheduler Address',
    s.[status] AS 'Scheduler State'
FROM sys.dm_exec_requests r
CROSS APPLY sys.dm_exec_sql_text(r.sql_handle) qt
INNER JOIN sys.dm_os_tasks t
ON r.task_address = t.task_address
INNER JOIN sys.dm_os_workers w
ON t.worker_address = w.worker_address
INNER JOIN sys.dm_os_schedulers s
ON w.scheduler_address = s.scheduler_address
WHERE r.session_id = 53
```

Figure 1-12 shows the information that the query returned on my test SQL Server. To keep the results readable, I only selected columns from the DMVs to show the relation between them.

Session ID	Type of Request	Query Text	Task Address	Task State	Worker Address	Worker State	Scheduler Address	Scheduler State
1 53	SELECT	SELECT r.res...	0x00000001732844E8	RUNNING	0x0000000178A36160	RUNNING	0x000000017F130040	VISIBLE ONLINE

Figure 1-12. Results of the query from Listing 1-1

In the results we can see that Session ID 53 made a SELECT query request. I did a cross apply with the `sys.dm_exec_sql_text` Dynamic Management Object to show the query text of the request. The request was mapped to a task, and the task began running. The task was then mapped to a worker thread that was then also in a running state. This meant that this query began being processed on a processor. The Scheduler Address column shows on which specific scheduler our worker thread was being run.

Wait Statistics

So far we have gone pretty deep into the different components that perform scheduling for SQL Server and how they are interconnected, but we haven't given a lot of attention to the topic of this book: Wait Statistics.

In the section about worker threads earlier in this chapter, I described the states a worker thread can be in while it is performing work on a scheduler. When a worker thread is performing its work, it goes through three different phases (or queues) in the scheduler process. Depending on the phase (or queue) a worker thread is in, it will get either the "RUNNING," "RUNNABLE," or "SUSPENDED" state. Figure 1-13 shows an abstract view of a scheduler with the three different phases.

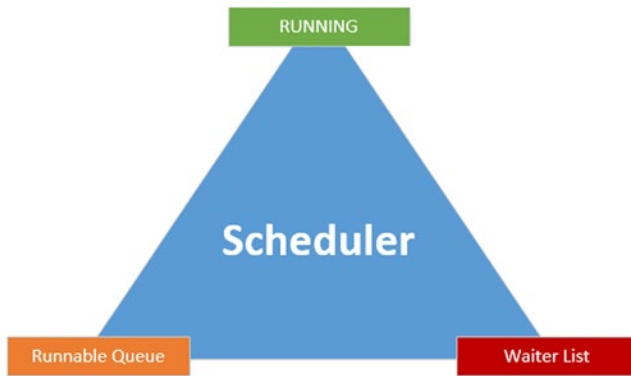


Figure 1-13. Scheduler and its phases and queues

When a worker thread gets access to a scheduler it will generally start in the Waiter List and get the "SUSPENDED" state. The *Waiter List* is an unordered list of worker threads that have the "SUSPENDED" state and are waiting for resources to become available. Those resources can be just about anything on the system, from data pages to a lock request. While a worker thread is in the Waiter List the SQLOS records the type of resource it needs to continue its work (the Wait Type) and the time it spends waiting before that specific resource becomes available, known as the *Resource Wait Time*.

Whenever a worker thread receives access to the resources it needs, it will move to the *Runnable Queue*, a first-in-first-out list of all the worker threads that have access to their resources and are ready to be run on the processor. The time a worker thread spends in the Runnable Queue is recorded as the *Signal Wait Time*.

The first worker thread in the Runnable Queue will move to the “RUNNING” phase, where it will receive processor time to perform its work. The time it spends on the processor is recorded as *CPU time*. In the meantime, the other worker threads in the Runnable Queue will move a spot higher in the list, and worker threads that have received their requested resources will move from the Waiter List into the Runnable Queue.

While a worker thread is in the “RUNNING” phase there are three scenarios that can happen:

- The worker thread needs additional resources; in this case it will move from the “RUNNING” phase to the Waiter List.
- The worker thread spends its quantum (fixed value of 4 milliseconds) and has to yield; the worker thread is moved to the bottom of the Runnable Queue.
- The worker thread is done with its work and will leave the scheduler.

Worker threads move through the three different phases all the time, and it is very common that one worker thread moves through them multiple times until its work is done.

Figure 1-14 will show you the scheduler view from Figure 1-13 combined with the different types of Wait Time and the flow of Worker Threads.




Figure 1-14. Scheduler view complete with Wait Times and worker thread flow

Knowing all the different lengths of time a request spends in one of the three different phases makes it possible to calculate the total request execution time, and also the total time a request had to wait for either processor time or resource time. Figure 1-15 shows the calculation of the total execution time and its different parts.

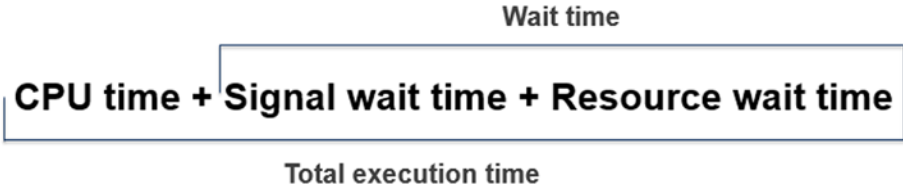


Figure 1-15. Request execution time calculation

Since there is a lot of terminology involved into the scheduling of worker threads in SQL Server, I would like to give you an example on how worker threads move through a scheduler.

Figure 1-16 will show you an abstract image of a scheduler like those we have already looked at, but this time I added requests that are being handled by that scheduler.

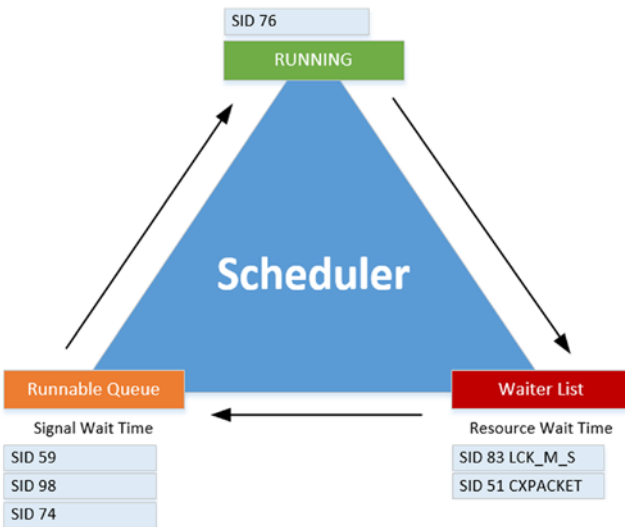


Figure 1-16. Scheduler with running requests

In this example we see that the request from **SID** (Session ID) **76** is currently being executed on the processor; this request will have the state “RUNNING.” There are two other requests, **SID 83** and **SID 51**, in the Waiter List waiting for their requested resources. The Wait Types they are waiting for are LCK_M_S and CXPACKET. I won’t go into detail here about these Wait Types since we will be covering both of them in Part II of this book. While these two sessions are in the Waiter List, SQL Server will be recording the time they spend there as Wait Time, and the Wait Type will be noted as the representation of the resource they are waiting on. If we were to query information about these two threads, they would both have the “SUSPENDED” state. **SID 59**, **SID 98**, and **SID 74** have their resources ready and are waiting in the Runnable Queue for **SID 76** to complete its work

on the processor. While they are waiting in the Runnable Queue, SQL Server records the time they spend there as the Signal Wait Time and adds that time to the total Wait Time. These three worker threads will have the status of “RUNNABLE.”

In Figure 1-17 we have moved a few milliseconds forward in time; notice how the scheduler and worker threads have moved through the different phases and queues.

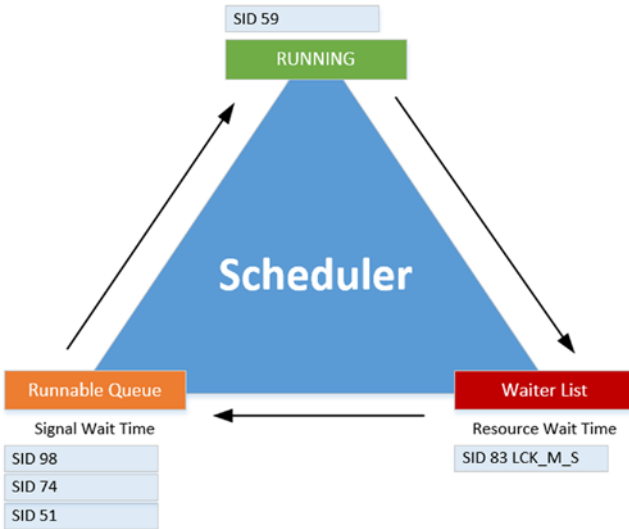


Figure 1-17. Scheduler a few milliseconds later

SID 76 completed its time on the processor; it didn’t need any additional resources to complete its request and thus left the scheduler. **SID 59** was the first worker thread in the Runnable Queue, and now that the processor is free it will move from the Runnable Queue to the processor, and its state will change from “RUNNABLE” to “RUNNING.” **SID 51** is done waiting on the CXPACKET Wait Type and moved from the Waiter List to the bottom of the Runnable Queue, changing its state from “SUSPENDED” to “RUNNABLE.”

Summary

In this chapter we took a look at the history of Wait Statistics throughout various versions of SQL Server. Even though the method of analyzing SQL Server performance using Wait Statistics is relatively new, Wait Statistics have been a part of the SQL Server engine for a very long time.

With the introduction of the SQLOS in SQL Server 2005 a lot changed in how SQL Server processed requests, introducing schedulers, worker threads, and tasks. All the information for the various parts are stored in Dynamic Management Views (DMVs) or Dynamic Management Functions (DMFs), which are easily queried and return a lot of information about the internals of SQL Server.

Using these DMVs, we can view the progress of requests while they are being handled by a SQL Server scheduler and learn if they are waiting for any specific resources. The resources the requests are waiting for and the time they spend waiting for those resources are recorded as Wait Statistics, which is the main topic of this book.

CHAPTER 2



Querying SQL Server Wait Statistics

With the introduction of Dynamic Management Views (DMVs) in SQL Server 2005, viewing and analyzing Wait Statistics has become a lot easier and less tedious. In SQL Server versions prior to SQL Server 2005, we were limited to the `DBCC SQLPERF('WAITSTATS')` command to view Wait Statistics. Presently there are a variety of DMVs that return Wait Statistics-related information, and in this chapter we will take a detailed look at three of the most useful DMVs: `sys.dm_os_wait_stats`, `sys.dm_os_waiting_tasks`, and `sys.dm_exec_requests`.

Viewing Wait Statistics information is not only limited to DMVs though. We can also use the Windows Performance Monitor, or Perfmon, to view Wait Statistics information. SQL Server 2008 introduced yet another option to view Wait Statistics, Extended Events. While Extended Events were pretty complicated to work with in SQL Server 2008, meaning you would have to write an entire Extended Event session in T-SQL, Microsoft has drastically improved Extended Events in SQL Server 2012, making them a lot more user friendly and easier to use.

Sys.dm_os_wait_stats

The `sys.dm_os_wait_stats` DMV is probably one of the most important DMVs regarding Wait Statistics. This DMV is the replacement for the `DBCC SQLPERF('WAITSTATS')` command you would have had to use before SQL Server 2005. All of the information the `DBCC SQLPERF('WAITSTATS')` command returned is included in the `sys.dm_os_wait_stats` DMV, plus a little bit more.

The `sys.dm_os_wait_stats` DMV shows the total amount of Wait Time for every Wait Type since the start (or restart) of your SQL Server. It is also cumulative, adding Wait Time to the different Wait Types, resulting in an ever-increasing total. Querying the `sys.dm_os_wait_stats` DMV will give you insight into what your SQL Server has been waiting on the most since the time it started or was restarted. This can be helpful if you are looking for that grand total of Wait Time for every Wait Type, but many times you are interested in the Wait Time for a specific time segment. In this case it is possible to reset

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-1-4842-1139-7_2](https://doi.org/10.1007/978-1-4842-1139-7_2)) contains supplementary material, which is available to authorized users.

the `sys.dm_os_wait_stats` DMV without having to restart your SQL Server by using the `DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR)` SQL command. This will reset all Wait Statistics information back to 0 again, meaning you will lose all information before the reset. In Chapter 3, “Building a Solid Baseline,” we will take a look at a method that does not completely reset the `sys.dm_os_wait_stats` DMV.

As with every DMV in SQL Server, we can run a query against the `sys.dm_os_wait_stats` DMV just like we would against a table, in this case `SELECT * FROM sys.dm_os_wait_stats`. The results of this query are shown in Figure 2-1.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
77	SOS_MEMORY_TOPLEVELBLOCKALLOCATOR	0	0	0	0
78	SOS_PHYS_PAGE_CACHE	0	0	0	0
79	ONDEMAND_TASK_QUEUE	0	0	0	0
80	LOGMGR_QUEUE	1697	238533	155	60
81	REQUEST_FOR_DEADLOCK_SEARCH	48	239854	4999	239854
82	CHECKPOINT_QUEUE	0	0	0	0
83	PARALLEL_BACKUP_QUEUE	0	0	0	0
84	DUMP_LOG_COORDINATOR_QUEUE	0	0	0	0

Figure 2-1. `sys.dm_os_wait_stats`

Following are the available columns in the `sys.dm_os_wait_stats` DMV, along with a description of what each column can tell you:

- `wait_type` returns the Wait Type. The `sys.dm_os_wait_stats` will always return one row for every Wait Type possible in that specific SQL Server version.
- `waiting_tasks_count` shows a total of how many times a worker thread had to wait for that specific Wait Type.
- `wait_time_ms` returns the total Wait Time in milliseconds (1/1000 of a second) for that specific Wait Type since the start of the SQL Server instance or a manual reset of the DMV. This is the time a worker thread has spent in the Waiter List in the “SUSPENDED” state. It also includes the time the worker thread spent in the Runnable Queue in the “RUNNABLE” state while waiting for the scheduler to grant it processor time.
- `max_wait_time_ms` shows the maximum Wait Time in milliseconds a worker thread waited on that specific Wait Type.
- `signal_wait_time_ms` tells us the amount of time in milliseconds the worker thread spent in Runnable Queue waiting for processor time. Signal Wait Times are unavoidable and normal in OLTP systems where a large number of queries are being processed, all of them requesting time on the processor. The Signal Wait Time is also an important metric for detecting CPU pressure. Generally speaking, as it depends on the hardware of your system, seeing Signal Wait Time metrics higher than 25% of the total Wait Time can indicate CPU pressure, because the worker threads are waiting for the processor to become available instead of using resources.

You may have noticed that the `sys.dm_os_wait_stats` DMV does not return a column for the Resource Wait Time. If we want to display the Resource Wait Time as an additional column we will need to calculate the value ourselves.

Listing 2-1 shows a query you could use to analyze the `sys.dm_os_wait_stats` DMV. Besides the regular columns, it will add two more columns for every Wait Type returned, the Resource Wait Time, and the average Wait Time.

Listing 2-1. `sys.dm_os_wait_stats` with additional information

```
SELECT
    wait_type AS 'Wait Type',
    waiting_tasks_count AS 'Waiting Tasks Count',
    (wait_time_ms - signal_wait_time_ms) AS 'Resource Wait Time',
    signal_wait_time_ms AS 'Signal Wait Time',
    wait_time_ms AS 'Total Wait Time',
    COALESCE(wait_time_ms / NULLIF(waiting_tasks_count,0), 0) AS 'Average Wait
Time'
FROM sys.dm_os_wait_stats;
```

This query will return results as shown in Figure 2-2.

	Wait Type	Waiting Tasks Count	Resource Wait Time	Signal Wait Time	Total Wait Time	Average Wait Time
46	TRAN_MARKLATCH_DT	0	0	0	0	0
47	LAZYWRITER_SLEEP	2506	2526986	512	2527498	1008
48	IO_COMPLETION	0	0	0	0	0
49	ASYNC_IO_COMPLETION	0	0	0	0	0
50	ASYNC_NETWORK_IO	10	142	23	165	16
51	SLEEP_BPOOL_FLUSH	0	0	0	0	0
52	CHKPT	0	0	0	0	0
53	SLEEP_DBSTARTUP	0	0	0	0	0

Figure 2-2. `sys.dm_os_wait_stats` expended with more wait information

Having both the number of occurrences of a specific Wait Type and the total Wait Time makes it possible to calculate an average Wait Time (represented by the Average Wait Time column in Figure 2-2) for that specific Wait Type by dividing the `wait_time_ms` value by the `waiting_tasks_count` value.

The `sys.dm_os_wait_stats` is a powerful DMV with which you can retrieve a lot of information about the different Wait Types. It is also the basis of the Wait Statistics Baseline methodology outlined in Chapter 3, “Building a Solid Baseline.”

Sys.dm_os_waiting_tasks

While the `sys.dm_os_wait_stats` DMV gives you the cumulative Wait Statistics information, the `sys.dm_os_waiting_tasks` DMV can give you information about what your SQL Server is waiting on right now! Querying this DMV will give you an overview of all the tasks that currently have worker threads waiting in the Waiter List or Runnable Queue for either resource or processor time.

Because the `sys.dm_os_waiting_tasks` DMV gives insight into what's waiting right now, it is usually the first DMV you query when using Wait Statistics to review the performance of your SQL Server instance. It also supplies some additional information for certain Wait Types that can be useful while troubleshooting.

Figure 2-3 shows the results of a `SELECT * FROM sys.dm_os_waiting_tasks` query.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address	blocking_task_address	blocking_sess
1	0x000000017f0168c8	24	0	63126	XE_DISPATCHER_WAIT	NULL	NULL	NULL
2	0x000000017f016cA8	1	0	56348	QDS_PERSIST_TASK_MAIN_LOOP_SLEEP	NULL	NULL	NULL
3	0x000000017f017088	3	0	5	SLEEP_TASK	NULL	NULL	NULL
4	0x000000017f017848	18	0	67149718	ONDEMAND_TASK_QUEUE	0x000000006790E820	NULL	NULL
5	0x000000017f017c28	7	0	3345	REQUEST_FOR_DEADLOCK_SEARCH	0x00000000682DE840	NULL	NULL
6	0x000000017BADA108	8	0	67145175	KSOURCE_WAKEUP	NULL	NULL	NULL
7	0x000000017BADA8C8	10	0	67148848	DISPATCHER_QUEUE_SEMAPHORE	NULL	NULL	NULL
8	0x00000001756C88C8	27	0	67145435	BROKER_TRANSMITTER	NULL	NULL	NULL
9	0x00000001756C8CA8	28	0	300	HADR_FILESTREAM_IOMGR_IOCOMPLE..	NULL	NULL	NULL

Figure 2-3. `sys.dm_os_waiting_tasks`

Following is a list of columns returned by the `sys.dm_os_waiting_tasks` DMV and a description of the information they return.

- `waiting_task_address` shows the address of the task that is currently waiting.
- `session_id` gives us the ID of the session that is associated with the specific task.
- `exec_context_id` will return the ID of the execution context. This value will only change from the default of 0 if the task is being performed using parallelism. This means that the task is being executed using multiple threads instead of a single (serial) thread.
- `wait_duration_ms` shows us the time in milliseconds that the task has been waiting. Just like in the `sys.dm_os_wait_stats` DMV, this time includes both the Resource Wait Time and the Signal Wait Time.
- `wait_type` returns the Wait Type the task is currently waiting on.
- `resource_address` returns memory address information about the resource we are currently waiting for. Not all Wait Types will log this memory address, so it will frequently be returned as NULL.
- `blocking_task_address` will return the address of the task that is currently blocking the waiting task. When the task is not being blocked by another task, this column will return NULL.
- `blocking_session_id` returns the session ID of the session that is currently blocking the task. Just like the `blocking_task_address`, this information is only included when this task is being currently blocked by another task. It will return NULL when there is no blocking or when the session information about the blocking task cannot be retrieved or identified. We will explain blocking and locking in Chapter 7 when we discuss lock Wait Types.

- `blocking_exec_context_id` is another column dedicated to information regarding blocking. In this case it will return the ID of the execution context. This will only return a result other than NULL when a task gets executed using parallelism and one of the threads is responsible for the block. The `blocking_exec_context_id` can then be used to identify which one of the threads is responsible for the block.
- The last column of the DMV, `resource_description`, will give additional information about the resource the task is waiting for. There aren't many Wait Types that will fill this column—most often parallelism, lock-, or latch-related Wait Types. It can be a very useful column, especially when analyzing lock- or latch-related Wait Types; in those cases we can pinpoint the database object (data page, row, table, etc.) whose availability we are waiting for. Some of the examples later in this book (most notably Chapter 7, “Lock Wait Types,” and Chapter 8, “Latch Wait Types”) will make use of this column to gather extra information about the resource we are waiting for.

Querying `sys.dm_os_waiting_tasks`

Because the `sys.dm_os_waiting_tasks` DMV returns a wealth of information, there are various ways to query it depending on what you want to analyze or troubleshoot.

One query that I see a lot on various forums on the Internet is the following:

```
SELECT * FROM sys.dm_os_waiting_tasks
WHERE session_id > 50;
```

This query will filter out all of the SQL Server internal session IDs and only returns waiting tasks that originate from user sessions. The results of the query on my test SQL Server are shown in Figure 2-4.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address	blocking_task_address
1	0x00000001763AB848	55	0	232	PAGEIOLATCH_SH	0x0000000004D82000	NULL

Figure 2-4. `sys.dm_os_waiting_tasks` where `session_id` is greater than 50

While the method of filtering out internal SQL Server processes will work fine for many Wait Types and improves readability, there are specific Wait Types that will not be returned when running this query.

One good example of this is the `THREADPOOL` Wait Type, which we will discuss in Chapter 4, “CPU-Related Wait Types.” This Wait Type can have a large negative impact on the performance of your SQL Server but will not be returned if you only query the user sessions. This can impact your analysis, because you are missing important facts.

Another reason to query the DMV without filtering on session IDs is that there is a big misconception about the relation of the session ID and whether or not the session ID is a user or internal session. While session IDs larger than 50 are generally considered to

be user sessions, there is no guarantee that a session ID larger than 50 actually *is* a user session. It is possible that there is a need for SQL Server to have more than 50 internal sessions, in which case there is a chance you will see internal sessions with a session ID higher than 50 that you can mistake for a user session.

I believe the best way to query the `sys.dm_os_waiting_tasks` DMV is by selecting everything and only applying a filter if you are looking for a specific Wait Type or session. This will return many more rows than filtering on session IDs larger than 50, as you can see in Figure 2-5, but it will show you the complete picture and minimizes the chance that you might miss important Wait Types. A good idea might be to sort on the `session_id` column to make the results a little bit more readable without losing sight of the internal sessions.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address
1	0x0000000180016CAB	25	0	51210	XE_DISPATCHER_WAIT	NULL
2	0x0000000180017088	1	0	53766	QDS_PERSIST_TASK_MAIN_LOOP_SLEEP	NULL
3	0x0000000180017468	3	0	8	SLEEP_TASK	NULL
4	0x0000000180017C28	16	0	6649871	ONDEMAND_TASK_QUEUE	0x000000067B5EF00
5	0x000000017CA06108	7	0	2914	REQUEST_FOR_DEADLOCK_SEARCH	0x00000006943E900
6	0x000000017CA064E8	8	0	6640314	KSOURCE_WAKEUP	NULL
7	0x000000017CA06CA8	9	0	2532386	DISPATCHER_QUEUE_SEMAPHORE	NULL
8	0x00000001763C5468	11	0	6646015	BROKER_TRANSMITTER	NULL
9	0x00000001763C5848	27	0	878	BROKER_TO_FLUSH	NULL
10	0x0000000180020CA8	24	0	4574	XE_TIMER_EVENT	NULL
11	0x0000000180021088	2	0	53859	QDS_CLEANUP_STALE_QUERIES_TASK...	NULL
12	0x0000000180021C28	18	0	2517493	CHECKPOINT_QUEUE	0x000000067DEF170
13	0x0000000180021848	6	0	242	LAZYWRITER_SLEEP	NULL
14	0x000000017CA089C8	10	0	2531827	WAIT_XTP_HOST_WAIT	NULL
15	0x000000017CA09468	14	0	50207	SP_SERVER_DIAGNOSTICS_SLEEP	0x0000000000000001
16	0x00000001763AA8C8	26	0	6640476	BROKER_EVENTHANDLER	NULL
17	0x00000001763AACAA8	28	0	6645949	BROKER_TRANSMITTER	NULL
18	0x00000001763AB088	29	0	456	HADR_FILESTREAM_IOMGR_IOCOMPLE...	NULL

Figure 2-5. `sys.dm_os_waiting_tasks`

Sys.dm_exec_requests

The `sys.dm_exec_requests` DMV returns information about all the requests that are currently getting processed by SQL Server.

Like the previous DMVs, we can query the `sys.dm_exec_requests` DMV with a simple `SELECT * FROM sys.dm_exec_requests`; to return everything that is currently executing. Figure 2-6 returns a small portion of the results on my test SQL Server.

	session_id	request_id	start_time	status	command	sql_handle	statement_start_offset
1	1	0	2015-02-23 23:50:47.800	background	UNKNOWN TOKEN	NULL	NULL
2	2	0	2015-02-23 23:50:47.813	background	UNKNOWN TOKEN	NULL	NULL
3	3	0	2015-02-23 23:50:47.813	background	UNKNOWN TOKEN	NULL	NULL
4	4	0	2015-02-23 23:50:47.990	background	LOG WRITER	NULL	NULL
5	5	0	2015-02-23 23:50:47.993	background	RECOVERY WRITER	NULL	NULL
6	6	0	2015-02-23 23:50:47.993	background	LAZY WRITER	NULL	NULL
7	7	0	2015-02-23 23:50:48.027	background	LOCK MONITOR	NULL	NULL

Figure 2-6. `sys.dm_exec_requests`

The `sys.dm_exec_requests` DMV returns a lot more columns than the `sys.dm_os_wait_stats` or `sys.dm_os_waiting_tasks` DMVs we discussed earlier. To keep things readable I will only describe the columns we will frequently use for Wait Statistics analysis. Here is the list:

- `session_id` returns the ID of the session this request is associated with.
- `start_time` shows the date and time when the request got created. This can be a different time and even date than the time and date on which you are querying the DMV, especially when long-running queries are being executed.
- `command` returns information about what kind of action the request is performing. The most common commands are query related, like `SELECT`, `INSERT`, `UPDATE`, and `DELETE`, but there are many more commands depending on what is being executed by the request.
- `sql_handle` gives us a hash value of the SQL text that is being executed in the request. Not all requests have an SQL handle, and generally you should only see a SQL handle if the request was initiated by a user session and an SQL query is involved. The SQL handle hash can be used as input for the Dynamic Management Function (DMF) `sys.dm_exec_sql_text` to retrieve the query that is being executed by the request.
- `plan_handle` returns the hash value of the execution plan. An execution plan will show you the operations that were performed by SQL Server when executing the query and is a great source of query-execution information. We can use `plan_handle` the same way as the `sql_handle`, but instead of returning the query it will return the execution plan of the query. We can use the hash value as input for the DMF `sys.dm_exec_query_plan` in order to return the execution plan of the query that is being executed by the request.
- `wait_type` returns the current Wait Type if the request is either “SUSPENDED” or “RUNNABLE.” The value will be `NULL` if the request is currently being processed.
- The `last_wait_type` column returns the last Wait Type the request encountered if it had to wait during its execution.
- `total_elapsed_time` column returns the total time, in milliseconds, it took to process the request from the moment it got created.

There are many more columns available in this DMV that all have their different uses. A complete description is available on the Microsoft MSDN page at <https://msdn.microsoft.com/en-us/library/ms177648.aspx>, and I encourage you to go through the article. The `sys.dm_exec_requests` DMV is a great tool in your DBA toolkit and is one of those DMVs you will use frequently for all kinds of purposes besides analyzing Wait Statistics.

Querying sys.dm_exec_requests

The `sys.dm_exec_requests` DMV is one of the DMVs that can give us access to query statements and corresponding execution plans by returning the query and plan handles. If you are interested in this information, and most of the time you probably will be, you need to pass the `sql_handle` and `plan_handle` to their DMFs so the hashes turn into something we humans can read and understand.

Listing 2-2 shows a query against the `sys.dm_exec_requests` DMV and also retrieves the query statements and execution plans. I am excluding the session ID I am executing the query on and ignoring session IDs lower than 50 so as to keep the results small, and because I know for a fact that I am interested only in user queries for this example.

Listing 2-2. Query `sys.dm_exec_request` and include query statement and plan

```
SELECT
  r.session_id AS 'Session ID',
  r.start_time AS 'Request Start',
  r.[status] AS 'Current State',
  r.[command] AS 'Request Command',
  t.[text] AS 'Query',
  p.query_plan AS 'Execution Plan'
FROM sys.dm_exec_requests r
OUTER APPLY sys.dm_exec_sql_text(r.sql_handle) AS t
OUTER APPLY sys.dm_exec_query_plan(r.plan_handle) p
WHERE r.session_id > 50 AND r.session_id <> @@SPID;
```

On my test system I got the results as shown in Figure 2-7.

	Session ID	Request Start	Current State	Request Command	Query	Execution Plan
1	55	2015-02-24 04:40:31.143	suspended	SELECT	SELECT * FROM Sales.SalesOrderDetail ORDER BY...	<ShowPlanXML x...

Figure 2-7. Results of Listing 2-2

Using the query in Listing 2-2 we can immediately see that session ID 55 is performing a SELECT query, as the query column shows the complete statement that is being executed. The Execution Plan column returns the execution plan in an XML format. The great part of the SQL Server Management Studio is that we can click the XML link that is returned to view the graphical execution plan, as shown in Figure 2-8.



Figure 2-8. Execution plan

Using the execution plan we can get some insight into how the query is getting executed by the SQL Server engine. We won't get into the details about execution plans in this book, but you will be using them frequently when you are optimizing query performance, so it is good to know how you can access them from the `sys.dm_exec_requests` DMV. A good place to start if you want to learn more about execution plans is Grant Fritchey's *Execution Plan Basics* at <https://www.simple-talk.com/sql/performance/execution-plan-basics/>.

Combining DMVs to Detect Waits Right Now

Now that we have taken a look at some of the most important DMVs for Wait Statistics analysis, let's go into an example of how you could use these DMVs to find out what is slowing down your SQL Server. Gathering this information will not solve your problems immediately, but it will give you a clue as to where to start looking for a solution.

Consider the following scenario: You are the database administrator (DBA) for a large company that uses a single database to store all its sales information. The database is running on a SQL Server 2014 instance, and every day a few hundred users query the database.

Normally everything is running fine—users can access the information they want quickly, and everyone who needs to work with the database is happy. Today, however, is not a good day for you as the DBA. The phone hasn't stopped ringing since 10 AM and some users are gathering at the door of your office with an angry look in their eyes—querying and inserting sales information is incredibly slow.

Since this book is about Wait Statistics, let's take a look at how we could analyze Wait Statistics information about the performance problem in the scenario.

We know the `sys.dm_os_wait_stats` DMV shows cumulative Wait Statistics information, so for this scenario it wouldn't be much help. A much better starting place would be the `sys.dm_os_waiting_tasks` DMV, since it will show us all the tasks that are waiting right now.

We run the following query against the `sys.dm_os_waiting_tasks` DMV:

```
SELECT * FROM sys.dm_os_waiting_tasks
ORDER BY session_id ASC;
```

While scrolling down to the bottom of the results, we see a number of user sessions with waiting tasks, as shown in Figure 2-9.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address
23	0x00000001763C5848	27	0	177	SLEEP_TASK	NULL
24	0x00000001763AAC8A	28	0	53986508	BROKER_TRANSMITTER	NULL
25	0x00000001763AB088	29	0	190	HADR_FILESTREAM_IOMGR_JOCOMPLETION	NULL
26	0x0000000174314108	52	0	194266	LCK_M_S	0x0000000163E12900
27	0x0000000176EF4CA8	55	0	148798	LCK_M_S	0x0000000176CC86C0
28	0x0000000173D59088	56	0	144617	LCK_M_S	0x0000000174C96940
29	0x0000000174314CA8	61	0	135854	LCK_M_S	0x0000000163E00680

Figure 2-9. Results of a query against the `sys.dm_os_waiting_tasks` DMV

We notice that the Wait Times for sessions 52, 55, 56, and 61 are pretty high, and we also notice that they are all waiting with a Wait Type LCK_M_S. Without going into too many details about this particular Wait Type, it is enough to know that this Wait Type is related to locking. Apparently those sessions are waiting to place a lock, which means they are probably being blocked by another process that has a lock on the same object. We can extract locking and blocking information from the `sys.dm_os_waiting_tasks` DMF by looking at the `blocking_` columns. For readability reasons I modified the query above to return only blocking information from the `sys.dm_os_waiting_tasks` DMV; Figure 2-10 shows those columns.

	session_id	wait_type	blocking_task_address	blocking_session_id	blocking_exec_context_id
22	26	BROKER_EVENTHANDLER	NULL	NULL	NULL
23	27	SLEEP_TASK	NULL	NULL	NULL
24	28	BROKER_TRANSMITTER	NULL	NULL	NULL
25	29	HADR_FILESTREAM_IOMGR_IOCOMPLETION	NULL	NULL	NULL
26	52	LCK_M_S	NULL	54	NULL
27	55	LCK_M_S	0x0000000174314108	52	NULL
28	56	LCK_M_S	0x0000000174314108	52	NULL
29	61	LCK_M_S	0x0000000174314108	52	NULL

Figure 2-10. Blocking information from the `sys.dm_os_waiting_tasks` DMV

From what we can see here, sessions 55, 56, and 61 are being blocked by session 52. Session 52, however, is being blocked by session ID 54. We don't see this session ID returned in the `sys.dm_os_waiting_tasks` DMV, which means the session is currently executing and isn't waiting on any resources.

Let's check another DMV, `sys.dm_exec_requests`, to get some information about session 54:

```
SELECT * FROM sys.dm_exec_requests
WHERE session_id = 54;
```

Figure 2-11 shows the results of this query.

session_id	request_id	start_time	status	command	sql_handle	statement_start_offset	statement_end_offset	plan_handle

Figure 2-11. Results of a query against `sys.dm_exec_requests`

Remember when I wrote the `sys.dm_exec_requests` DMV return information about requests currently being processed? Apparently session ID 54 doesn't have an outstanding request, since no information is being returned.

If we want to find out more information about this session we can use the `sys.dm_exec_sessions` DMV we discussed in Chapter 1 by executing the following query:

```
SELECT
  session_id,
  [status],
  [host_name],
```

```
[program_name],
login_name,
is_user_process,
open_transaction_count
FROM sys.dm_exec_sessions
WHERE session_id = 54;
```

This query returns the results shown in Figure 2-12.

	session_id	status	host_name	program_name	login_name	is_user_process	open_transaction_count
1	54	sleeping	EVDL-SQL2014-01	Microsoft SQL Server Ma...	EVDL-SQL2014-01\Administrator	1	1

Figure 2-12. Results from `sys.dm_exec_sessions`

We can assume that session ID doesn't currently have any running requests since its status is "sleeping," which is why the query against the `sys.dm_exec_requests` didn't return any information. If we look at the `program_name` column we can see that this session was initiated from the Microsoft SQL Server Management Studio program by the `EVDL-SQL2014-01\Administrator` user.

I included the `is_user_process` column to make sure it is a user session, the `open_transaction_count` column shows us that this user session has an open transaction.

We now know enough information to take corrective actions. We know who the user is who is blocking our other tasks, and we can decide to give him a call about what he is currently performing, or we can choose to end his session. Ending a user session by using the `KILL [session_id]` command should always be your last resort, because we could be interrupting something important. Ending a session with the `KILL` command will result in a rollback of the running transaction, undoing all the changes it performed, which can take a long time to complete. In this case, I accept the risk of a rollback and will end the session myself:

```
KILL 54;
```

Immediately after we kill session ID 54 users report that their queries are running again. If we query the `sys.dm_os_waiting_tasks` DMV to give us information about those session IDs nothing gets returned, meaning they are no longer being blocked.

Hopefully this example has given you insight into how you can use the various DMVs available in SQL Server to gather information about tasks that are currently waiting. In this case the example consisted of a transaction that was blocking other queries, and we decided to kill the user session that was responsible for the blocking lock. In many situations the solution isn't this relatively simple, but the method of gathering Wait Statistics information to drill down to the bottom of the problem can be used in almost every performance-related incident.

As I noted in the beginning of this section, just looking at the Wait Statistics alone will not, in most cases, solve a performance problem, but it is a good starting point from which is to begin your investigation. To create a complete picture about the performance of your system we will often combine the Wait Statistics information with other metrics, like those from the Windows Performance Monitor, other DMVs, or vendor-specific information (like storage metrics).

Figure 2-13 shows a flowchart of how you could use Wait Statistics information to analyze a performance problem.

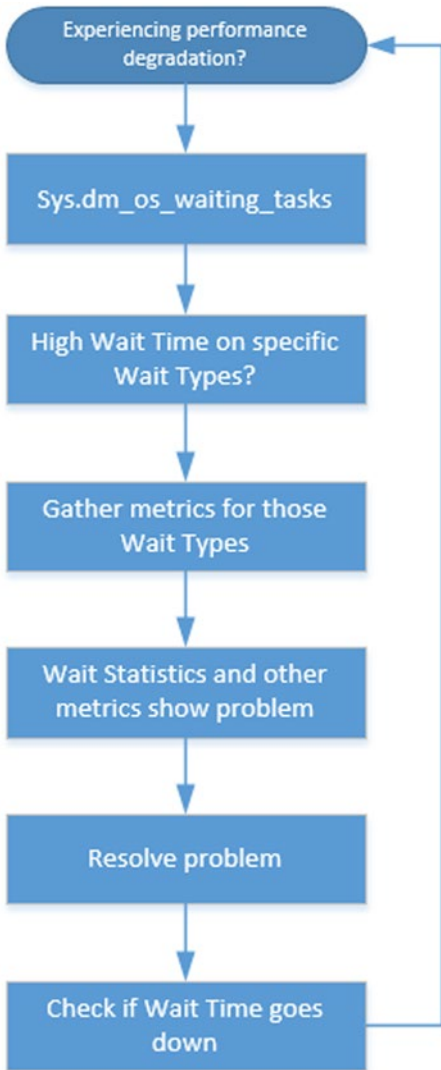


Figure 2-13. *Wait Statistics flowchart*

We will expand upon this flowchart in Chapter 3, where we will introduce baselines to the Wait Statistics analysis method.

Viewing Wait Statistics Using Perfmon

One of the most important tools for accessing the extra metrics we need when analyzing Wait Statistics is the Windows Performance Monitor, or Perfmon. Perfmon is available on every Windows Operating System and contains counters for just about every part of the system, including SQL Server–related performance counters. You can start Perfmon by executing the **perfmon** command from either a Windows Run dialog or the command line.

In addition to giving us information about the performance of our system, Perfmon can also be used to view Wait Statistics. You can view these counters under the `SQLServer:Wait Statistics` category when adding counters inside the Perfmon application, as shown in Figure 2-14.

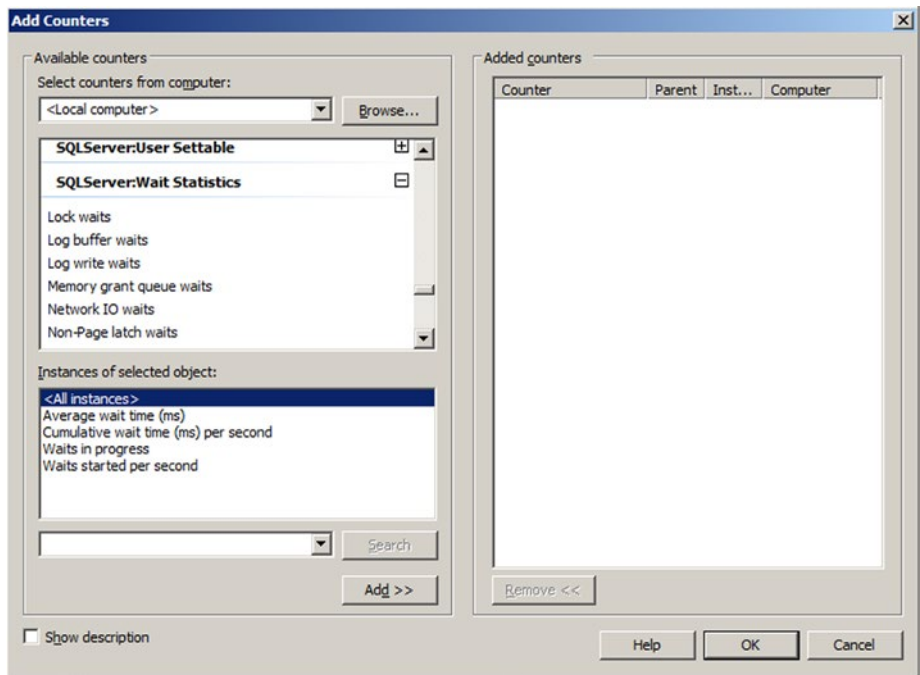


Figure 2-14. Wait Statistics counters inside Perfmon

One thing you’ll notice in Figure 2-14 is that the Wait Statistics inside Perfmon are grouped inside categories. We won’t find information about specific Wait Types here, so if we want to use Perfmon to analyze Wait Statistics we should have a general idea of what category a specific Wait Type belongs to in Perfmon. It is able to display an average wait time, cumulative wait time, the current total number of waits, and the amount of new waits started for every Wait Statistics category. If we are interested in a higher-level view—for instance, we want to know how many tasks are waiting for lock-related Wait Types—we can use Perfmon to give us that information. If we want to have more detail

about specific Wait Type information, we should use the `sys.dm_os_wait_stats` or `sys.dm_os_waiting_tasks` DMVs we discussed earlier.

One nice feature of Perfmon is that it can convert the measurements directly into graphs, giving us a more visual way to look at the information without having to create the graphs ourselves. Figure 2-15 is an example of a graph where we are showing the Average Wait Time and Waits started per second for the Lock waits category.




Figure 2-15. Perfmon graph showing Wait Statistics information

During this book we will use Perfmon a lot for analyzing metrics related to specific Wait Types, like CPU time, disk latency, and memory usage. We will not make much use of the Wait Statistics counters inside Perlman, because I believe the SQL Server DMVs are better suited for this since they supply the level of detail needed for a complete analysis.

Capturing Wait Statistics Using Extended Events

Most of the Wait Statistics information in SQL Server is recorded cumulatively, and because so many internal processes also generate Wait Statistics it can be difficult to detect what impact a single query has on them. This is where Extended Events come in; with Extended Events it is possible to capture the exact Wait Times a query encountered and on what Wait Types it had to wait. This information can help us analyze those queries that have a large impact on our system and possibly optimize them so their impact becomes smaller. Or we could capture queries that encounter a specific Wait Type while executing.

Extended Events were introduced in SQL Server 2008 and are, more or less, a replacement for the SQL Server Profiler. Microsoft has announced the deprecation of the SQL Server Profiler and advises us to move to Extended Events. Extended Events are much more powerful than the SQL Server Profiler, and the number of events we can

capture with Extended Events keeps growing with every release of SQL Server, while the amount of events in the SQL Server Profiler remains the same.

Extended Events have a reputation of being difficult to work with, and, while this was especially true in SQL Server 2008 when they were first introduced, working with Extended Events became a lot easier in SQL Server 2012 when it became possible to create Extended Event sessions using the GUI.

There are many different wait-related events available when working with Extended Events. We can view these events by running a query against the `sys.dm_xe_map_values` DMV, which holds all the different Extended Events event types:

```
SELECT *
FROM sys.dm_xe_map_values
WHERE name = 'wait_types';
```

Figure 2-16 shows a small part of the results of this query.

	name	object_package_guid	map_key	map_value
1	wait_types	BD97CC63-3F38-4922-AA93-607BD12E78B2	862	__indexMUTEX_HADR_DATABASE_WAIT_FOR_RESTART
2	wait_types	BD97CC63-3F38-4922-AA93-607BD12E78B2	864	__indexMUTEX_HADR_RECOVERY_WAIT_FOR_CONNECTION
3	wait_types	BD97CC63-3F38-4922-AA93-607BD12E78B2	861	__indexMUTEX_HADR_RECOVERY_WAIT_FOR_UNDO
4	wait_types	BD97CC63-3F38-4922-AA93-607BD12E78B2	387	ABR
5	wait_types	BD97CC63-3F38-4922-AA93-607BD12E78B2	751	ALL_COMPONENTS_INITIALIZED
6	wait_types	BD97CC63-3F38-4922-AA93-607BD12E78B2	755	AM_INDBUILD_ALLOCATION
7	wait_types	BD97CC63-3F38-4922-AA93-607BD12E78B2	850	AM_SCHEMAMGR_UNSHARED_CACHE
8	wait_types	BD97CC63-3F38-4922-AA93-607BD12E78B2	357	ASSEMBLY_LOAD
9	wait_types	BD97CC63-3F38-4922-AA93-607BD12E78B2	116	ASYNC_DISKPOOL_LOCK
10	wait_types	BD97CC63-3F38-4922-AA93-607BD12E78B2	98	ASYNC_IO_COMPLETION
11	wait_types	BD97CC63-3F38-4922-AA93-607BD12E78B2	904	ASYNC_OP_COMPLETION

Figure 2-16. Results of `sys.dm_xe_map_values`

In total there are about 931 different Wait Statistics–related events. These events do not map one-on-one against the different Wait Types, and, as a matter of fact, in some cases the names of the Wait Types do not match those of the events, even though they have the same meaning. An example of this is the `ASYNC_NETWORK_IO` Wait Type, which is named `NETWORK_IO` by Extended Events. Jonathan Kehayias wrote a blog post at [SQLskills.com](https://www.sqlskills.com/mapping-some-of-the-Wait-Types-to-Extended-Events/) mapping some of the Wait Types to Extended Events; you can take a look at it here: <https://www.sqlskills.com/blogs/jonathan/mapping-wait-types-in-dm-os-wait-stats-to-extended-events/>.

While we won't go into details about Extended Events in this book, I would like to show you how you can use them to capture Wait Statistics–related information using the Extended Events GUI and T-SQL.

Capture Wait Statistics Information for a Specific Query

Let's take a look at how we can configure an Extended Event session to capture Wait Statistics information for a specific query. We will set a filter on a session ID that will execute the query, then execute the query we want to analyze.

The first thing we are going to do is open up the SQL Server Management Studio and connect to a SQL Server instance. Keep in mind that a GUI for Extended Events was added in SQL Server 2012, so if you plan on following the steps here you will need a SQL Server 2012 or higher SQL Server instance.

Once we are connected, we open the Management folder and then choose the Extended Events option. We right-click the Sessions folder and select the option New Session, as shown in Figure 2-17.

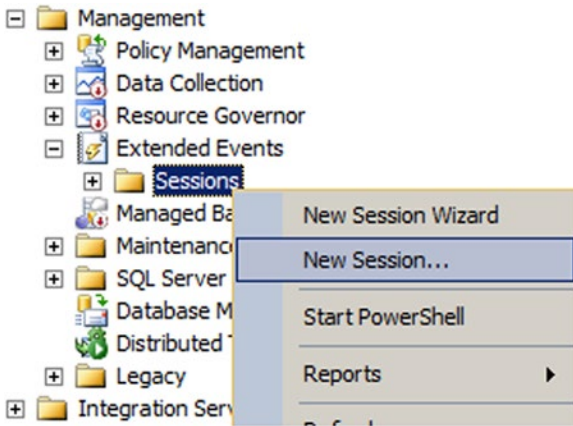


Figure 2-17. Adding new Extended Event session

The New Session dialog will appear where we can enter a name for this Extended Event session and set some additional options. We will ignore those options for now, and just fill in the name of the Extended Event session, as shown in Figure 2-18.

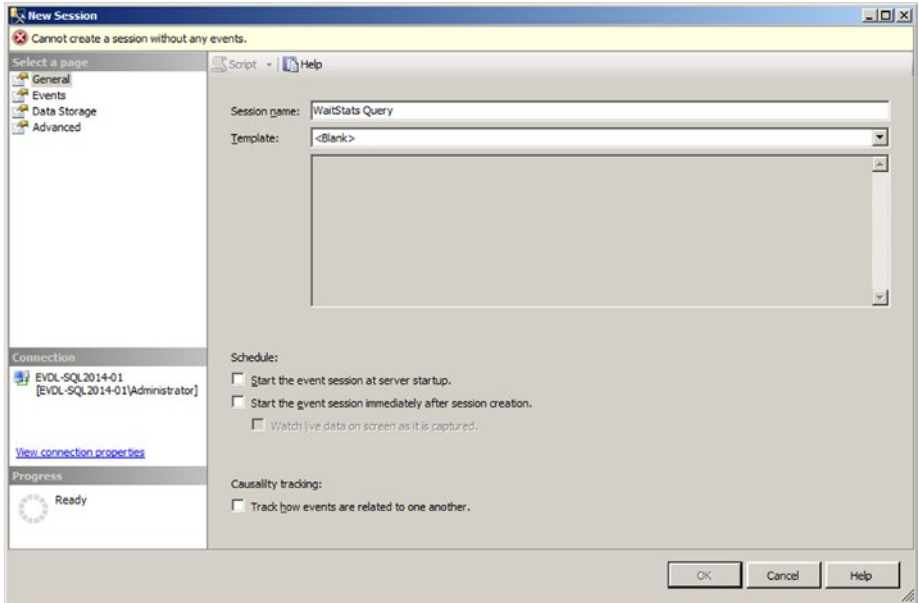


Figure 2-18. *Configuring Wait Statistics Extended Event session*

The next step is configuring which events this Extended Event session needs to monitor, which we can do by selecting the Events page in the New Session dialog.

Since we are interested in Wait Statistics information, I searched for the `wait_info` event in the Events Library and added it to the Selected Events box, as shown in Figure 2-19.

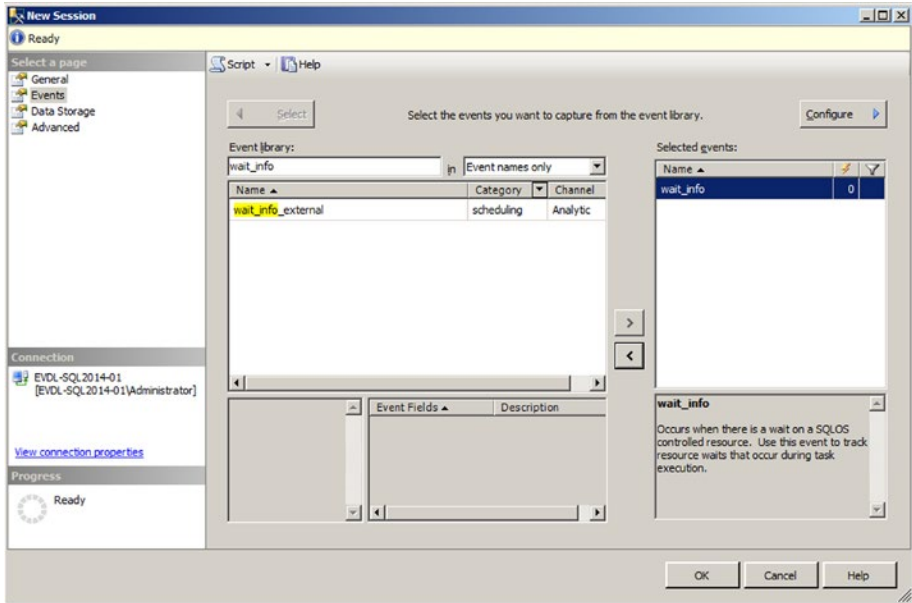


Figure 2-19. Selecting an event to monitor

If we were to save this Extended Event session now we would capture information for every task that has to wait for a resource. Since we are interested in the Wait Statistics associated with a specific query, we will configure a filter to only return Wait Statistics information for a specific session. To do this we can click the Configure button in the New Session dialog, which will open a new section where we can select Global Fields, which will record extra information when a `wait_info` event is triggered. In this case, I checked the `sql_text` global field, as shown in Figure 2-20, so we can view the actual query when an event is captured.

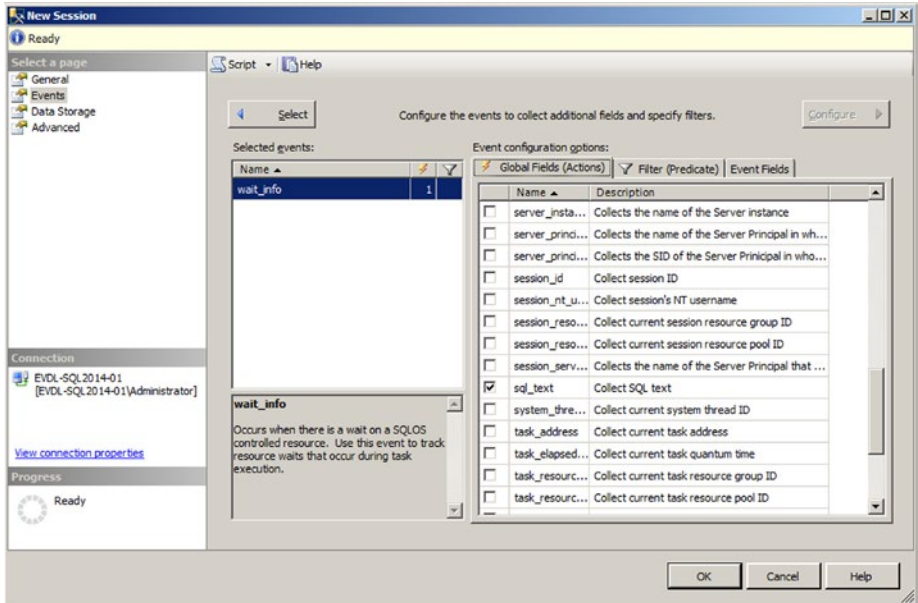


Figure 2-20. Setting the `sql_text` global field

Next up is the Filter (Predicate) tab. Here we will set a filter that will only capture events from a specific session ID. We can do this by clicking inside the Field box and selecting the `sqlserver.session_id` field, then setting the Value to the session ID we want to monitor. In this case I configured the filter to only capture events for session ID 53, as shown in Figure 2-21.

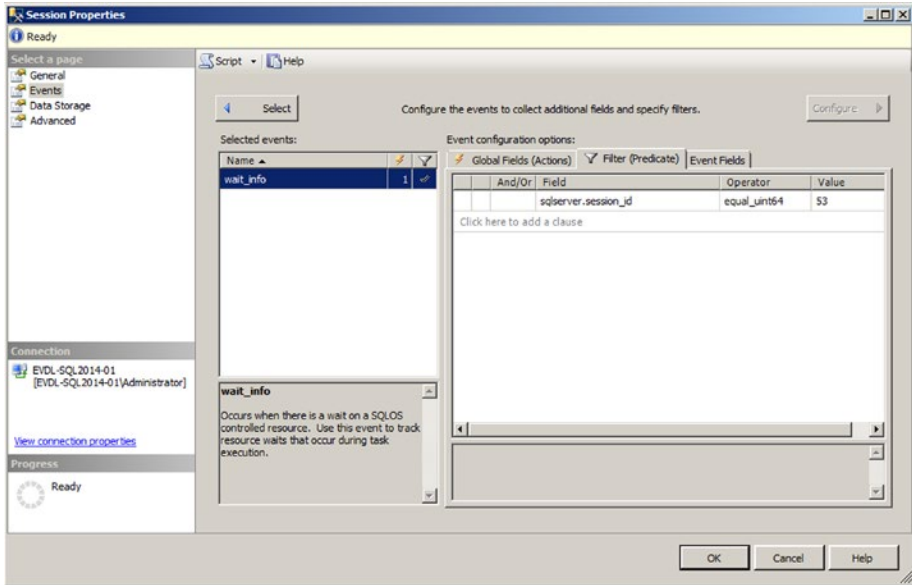


Figure 2-21. Setting an events filter

That’s all we need to configure for now, so we can click OK to close this dialog and save the Extended Event session.

By default the Extended Event session will not be automatically started after it is created. To do this we have to open up the Sessions folder again by navigating to the Management - Extended Events folder in SQL Server Management Studio. We right-click the Extended Event session we just created and select the Start Session option, as shown in Figure 2-22.

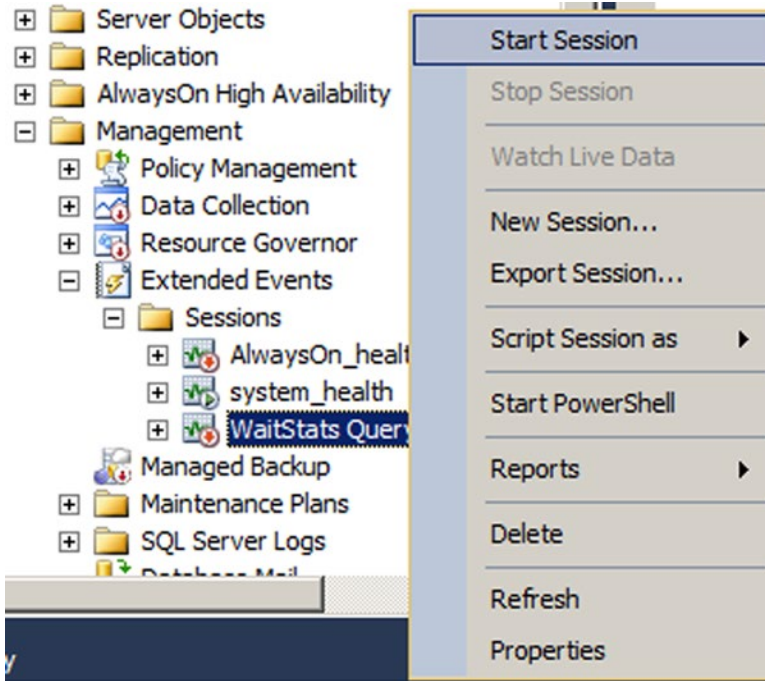


Figure 2-22. Start Extended Event session

After we have started the Extended Events session it will begin collecting information. We can view this information as it gets gathered by selecting the Watch Live Data option. This will open up a new tab in SQL Server Management Studio where we can watch the Extended Event session. Viewing live Extended Event data takes a little overhead, but this is far lower than the overhead of using SQL Profiler. If you are worried about the overhead of viewing live data, you could choose to write the Extended Event session to an event file by adding a file location inside the Data Storage page of your Extended Event session.

For this example I executed a simple query against the AdventureWorks2014 database to return everything in the Person.Person table, as follows:

```
SELECT *
FROM Person.Person;
```

I pay close attention when setting the filter in the Extended Event session to match the session ID of the tab in SQL Server Management Studio where I am executing the query. The session ID can be found when looking at the number between parentheses inside the tab. The Extended Events Live Data tab returned the information shown in Figure 2-23.

The screenshot shows a window titled 'SQLQuery25.sql - EVD... Administrator (53)*'. It displays '20917 Events' in a table with columns 'name' and 'timestamp'. Two rows are visible, both with 'wait_info' as the name and timestamps from 2015-03-02. Below this, the details for the selected event are shown in a table with columns 'Field' and 'Value'.

Event: wait_info (2015-03-02 11:46:58.2681171)	
Field	Value
duration	0
opcode	Begin
signal_duration	0
sql_text	SELECT * FROM Person.Person
wait_type	NETWORK_IO

Figure 2-23. Live Wait Statistics information from an Extended Event session

As you can see in Figure 2-23, our request encountered an `NETWORK_IO` Wait Type. This is one of those examples where the wait name in Extended Events doesn't match the one in the Wait Statistics DMVs. The `NETWORK_IO` wait name is the same wait as the `ASYNC_NETWORK_IO` Wait Type the SQLOS uses. We can view the query we executed in the `sql_text` field.

There are many more global fields you can include in the Extended Events session that might be interesting to capture, like the Execution Plan handle or the Task Execution Time. All of these global fields will give you additional information that is shown when Extended Event session information is returned, giving you an impressive amount of detail.

If for some reason you do not want to use the GUI to create and run an Extended Event session or you are running SQL Server 2008, you can use T-SQL to create and configure one. To create the same Extended Event session as we did using the GUI, you can execute the query seen in Listing 2-3.

Listing 2-3. Create Wait Statistics Extended Event session

```
CREATE EVENT SESSION [WaitStats Query] ON SERVER
ADD EVENT sqllos.wait_info
(
  ACTION(sqlserver.sql_text)
  WHERE ([sqlserver].[session_id]=(53))
)
```

```

ADD TARGET package0.event_file
(
    SET filename = N'E:\Data\WaitStats_XE.xel', metadatafile = N'E:\Data\
WaitStats_XE.xem'
);

```

We included the metadata file in the script above by setting the `metadatafile` parameter. If you are running SQL Server 2012 or higher this is no longer required.

The easiest way to log the Extended Event session is by saving it to a file; in this case my filename is `E:\Data\WaitStats_XE.xel` (SQL Server will add a unique numeric identifier to the filename, in this case the actual filename is `WaitStats_XE_0_130702270937280000.xel`). I also included the filter on session ID 53 to capture Wait Statistics generated by that session.

The next thing we want to do is start the Extended Event session, which we can do by executing the `ALTER EVENT SESSION` command:

```
ALTER EVENT SESSION "WaitStats Query" ON SERVER STATE = start;
```

We then execute the same query as we did in the Extended Events GUI example under the session ID we are filtering on. After letting the Extended Event session run for a little while we can stop it using the `ALTER EVENT SESSION` command:

```
ALTER EVENT SESSION "WaitStats Query" ON SERVER STATE = stop;
```

Now that we have stopped the Extended Event session, we need to import the information in the file (as XML) into a table so we can actually see what the session captured; we do this using the `sys.fn_xe_file_target_read_file` function. We can then parse the XML information to return the results in a more readable format. The query in Listing 2-4 can be used to read an Extended Events file, import it into a temporary table, and return the results as rows.

Listing 2-4. Return Extend Event File as Rows

```

-- Check if temp table is present
-- Drop if exist
IF OBJECT_ID('tempdb..#XE_Data') IS NOT NULL
DROP TABLE #XE_Data

-- Create temp table to hold raw XE data
CREATE TABLE #XE_Data
(
    XE_Data XML
);
GO

```



```

-- Write contents of the XE file
-- into our table
INSERT INTO #XE_Data
(
    XE_Data
)
SELECT
    CAST (event_data AS XML)
FROM sys.fn_xe_file_target_read_file
(
    'E:\Data\WaitStats_XE_O_130702270937280000.xel',
    'E:\Data\WaitStats_XE_O_130702270940210000.xem',
    null,
    null
);
GO

-- Query information from our temp table
SELECT
    XE_Data.value ('(/event/@timestamp)[1]', 'DATETIME') AS 'Date/Time',
    XE_Data.value ('(/event/data[@name='opcode']/text)[1]', 'VARCHAR(100)')
AS 'Operation',
    XE_Data.value ('(/event/data[@name='wait_type']/text)[1]',
'VARCHAR(100)') AS 'Wait Type',
    XE_Data.value ('(/event/data[@name='duration']/value)[1]', 'BIGINT') AS
'Wait Time',
    XE_Data.value ('(/event/data[@name='signal_duration']/value)[1]',
'BIGINT') AS 'Signal Wait Time',
    XE_Data.value ('(/event/action[@name='sql_text']/value)[1]',
'VARCHAR(100)') AS 'Query'
FROM #XE_Data
ORDER BY 'Date/Time' ASC
;

```

The result of the query in Listing 2-4 can be seen in Figure 2-24.

	Date/Time	Operation	Wait Type	Wait Time	Signal Wait Time	Query
1	2015-03-04 18:11:44.173	Begin	SOS_SCHEDULER_YIELD	0	0	SELECT * from
2	2015-03-04 18:11:44.173	End	SOS_SCHEDULER_YIELD	0	0	SELECT * from
3	2015-03-04 18:11:44.173	Begin	PAGEIOLATCH_SH	0	0	SELECT * from
4	2015-03-04 18:11:44.180	End	PAGEIOLATCH_SH	5	0	SELECT * from
5	2015-03-04 18:11:48.300	Begin	WRITELOG	0	0	ALTER EVENT
6	2015-03-04 18:11:48.300	End	WRITELOG	0	0	ALTER EVENT

Figure 2-24. Results of the query in Listing 2-4

Most of the columns speak for themselves in terms of the row data they return. Two columns that deserve some extra explanation are the `Operation` and `Wait Time` columns. The `Operation` column will show you the beginning or the end of the wait event. The `Wait Time` column will return the `Wait Time` in milliseconds, but it will only be recorded at the end of an operation.

Summary

In this chapter we reviewed the various ways we can access information about Wait Statistics. We took an in-depth look at some of the most important DMVs regarding Wait Statistics: `sys.dm_os_wait_stats`, `sys.dm_os_waiting_tasks`, and `sys.dm_exec_requests`. I described their functions and the data they returned, and gave you some example queries you can use against those DMVs. We also went through an example scenario where we combined some of the DMVs to analyze what was slowing down the SQL Server in the example. The steps shown in the example are a good way to analyze performance problems on your system when they are occurring. Briefly, we looked at the Windows Performance Monitor, or Perfmon, and how you can access Wait Statistics information from inside it. After that we took a good look at Extended Events and how you can use them to capture wait-related information for specific queries or sessions using the Extended Events GUI or T-SQL.

CHAPTER 3



Building a Solid Baseline

In the previous chapter we spent a great deal of time describing and using the various methods of accessing Wait Statistics information. Most of those methods focused on using that information for detecting performance problems that are presently occurring. While it is possible to find the exact cause of the performance problem using these real-time methods, it requires a deep knowledge of the various Wait Types and—most important—experience in the performance of your SQL Server. If you are managing only one SQL Server instance you can get yourself familiar with the way it reacts under different circumstances relatively quickly. If you are managing hundreds of SQL Server instances, getting yourself familiar with the way they perform is impossible. Because SQL Server Wait Statistics are largely based on the workload of your SQL Server instance, no two SQL Server instances will have the same Wait Times for the same Wait Types. This makes detecting possible problems difficult, because we can't just say “because the CXPACKET Wait Type has a Wait Time of 20,000 milliseconds we are having a problem.” It all depends on the configuration and workload of your system. One SQL Server instance can have 20,000 milliseconds (20 seconds) of Wait Time every minute spent on the CXPACKET Wait Type and experience no performance problems, while another instance has 1,000 milliseconds of Wait Time and users are constantly complaining about performance.

If we want to perform an in-depth analysis of Wait Statistics, or any performance-related data, we need a method of collecting performance-related metrics and giving them meaning. Just detecting that you spend 10,000 milliseconds waiting for resources doesn't mean anything, since we do not know if it caused performance problems or not. Yes, maybe your users are complaining that performance is horrible while you notice that 10,000 millisecond wait, but there is no way to be sure if that wait is actually causing the performance problems your users are experiencing. In those cases we frequently take a guess and just assume that the specific resource wait is the source of the performance problem. I learned, after talking to many DBAs across the world, that we DBAs do not like to take guesses at what's slowing down our SQL Server instances. We want to be certain that the source of the performance problem actually is that resource wait. This is where baselines can help.

Baselines will help you give meaning to performance metrics by providing you with a definition of what the normal situation is on your system. Without a baseline of our system we have no idea if it is running optimally or terribly slow. Baselines are incredibly important; in fact, they are so important that I decided to write a complete chapter about

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-1-4842-1139-7_3](https://doi.org/10.1007/978-1-4842-1139-7_3)) contains supplementary material, which is available to authorized users.

them in this book. Without a solid baseline your measurements mean nothing! Even though this book focuses on Wait Statistics, baselines can be used for every performance-related metric you can capture on your system, giving you a valuable method of performance analysis.

What Are Baselines?

If we look up the word “baseline” in the Oxford dictionary, we would get the following definition: “A minimum or starting point used for comparisons.” This sentence captures the essence of a baseline perfectly by using the important words “starting point” and “comparisons.” A baseline will generally be the starting point, or in our case, measurement, that we will compare later measurements against. Ideally, we will capture our baseline measurement in a normal or standard situation. If we perform the same measurement again at a later point in time, we can compare that measurement against the baseline. If our measurements are not the same during the comparison, something might have changed.

Even if you do not yet use baseline comparisons when analyzing performance, you are still working with baselines constantly whether you realize it or not. For instance, if you receive a salary every first day of the month, that would be your normal situation or, in the context of this chapter, your baseline. If for some reason you didn’t receive your salary on the first day of the month, you would notice a deviation compared to the baseline. This might be a reason to investigate why you didn’t receive your salary on time. Maybe the day changed from the first to the fifth of the month or, in the worst case, the company you work for can’t pay your salary anymore. In any of these cases there are actions we can take: either accept the change, and by doing so create a new baseline, or revert the situation back to the baseline state again. Figure 3-1 shows this process.




Figure 3-1. Changes impact baseline

Defining and maintaining baselines is an iterative process. Every change occurring on your system can impact your baseline. A new release of the application you are using might change a number of queries, or your company might grant a new department access to the database, increasing the number of connections. In both these examples we would need to make adjustments to the baseline, since the normal situation would have changed.

All of this adjusting and measuring baselines with every change to the system sounds like a lot of work, and sometimes it is. But believe me—the benefits of having a baseline far outweigh the costs. Baselines will help you detect problems far faster than just looking at a single measurement, and in the case of Wait Statistics it is the only way to find a reliable, definitive answer to your problem. Let's use a more technical example to illustrate this using DBA Jim.

Jim maintains a SQL Server instance that hosts a single user database. This database is used by every sales person in the company and records every financial transaction between the company and its customers. Users can access the database through Application X. Application X is currently running version 2.4 and is very stable. Performance is good, users are happy, and the money keeps rolling in. Sounds great, right? One day a consultant walks in and wants to upgrade Application X to the brand new 3.0 version. The update to version 3.0 was a breeze and completed without any problems, and all the users love the new features.

Two days later the phone is ringing, Chris's manager just received word from the sales team that the performance in version 3.0 is horrible and he demands it get resolved right away.

Thankfully, Jim knew the importance of a baseline, and he created one before the upgrade to version 3.0. Using the version 2.4 baseline, Jim compares the measurements in the baseline to the measurements done in version 3.0 and immediately spots a large difference in the lock Wait Time measurements. Since other measurements remain more or less the same compared to the 2.4 baseline, Jim focusses on long-running locks and identifies an update query that is locking a table. He rolls back the query, and the situation returns to normal. He then contacts the application's vendor and learns this behavior was a result of a bug in the software.

Now this example might sound a little far-fetched, but it is actually a simplified version of the method I use almost every day when measuring the impact of changes or analyzing performance problems. If Jim didn't have a baseline of the lock Wait Time and just queried the lock Wait Time after the change to version 3.0, he wouldn't know that the Wait Time had increased, since he had nothing to compare against. He might have chosen to look at other metrics instead of the lock Wait Time and would have wasted valuable time and money.

The message here is simple: baselines will help you detect abnormal situations and resolve performance problems faster!

Visualizing Your Baselines

Baselines are frequently visualized in the form of graphs. The big advantage of turning your baseline measurements into graphs is that a graph can make it easier for you to detect those measurements that have the highest increase or decrease compared to your baseline. Also, visualizing your data might help you get your point across easier if you have to prove that a specific configuration is impacting performance. An example: you need to convince your storage administrator that the change in the storage configuration has impacted your performance. If you can hand him a graph that shows the normal behavior compared to the behavior after the change, he might be more inclined to help. The graph in Figure 3-2 shows the baseline measurements compared to the measurements done at a later time.




Figure 3-2. Example of a baseline graph

As you can see in Figure 3-2, you would be able to identify the potential problem very quickly. Apparently, between 08:00 and 12:00 the number of transactions per second is higher than the normal situation, and it might be worth taking the time to investigate.

Baselines Types and Statistics

We will frequently use different types of baselines depending on the information we are interested in. There usually isn't one single baseline to fit all our needs, especially not when you are using baselines for performance-troubleshooting purposes. For instance, we can create a baseline for every single Wait Type, or we can choose to create only baselines for Wait Types that impact our system the most. We can also choose to create a baseline for specific days or time segments, like business hours, and create another for after business hours.

Next to selecting or limiting the measurements we want to have baselines for, we also have to make choices on how we calculate our baselines. These choices involve some math and usually require calculating averages. In many cases our baseline consists of an average of many data points, depending on how many measurements you performed. If you collect measurements for a long period of time and calculate an average value from those measurements, you can create a more reliable baseline than you can when you only have one day's worth of measurements. Creating a baseline based on averages also has its disadvantages. The most important one is that averages are heavily influenced by skewed data. Without going too deep into statistical details, skewed data means that there are very high or very low values that impact your average. Say, for instance, that a group of students took an exam and we wanted to see how the group performed by calculating the average result of the exam (the students are rated between 1 and 10, 1 being very poor and 10 being excellent; a 6 or higher is required to pass the exam). Figure 3-3 shows the exam results in a graph.




Figure 3-3. Exam results

As you can see in Figure 3-3, only four students scored higher than the required 6 to pass the exam. The rest of the group scored way below the requirement. However, if we look at the average performance of the group, they actually didn't do too badly by scoring a 6. We could conclude the group is performing well enough by getting an average score of 6, but then we would miss that the majority performed badly and only four students actually passed the exam. This information is important to keep in mind when you are dealing with average baselines. If you see a spike in your average baseline it is always something you should investigate, because it impacts your baseline.

There are statistical methods available to deal with skewed data and averages, one being the trimmed (or truncated) mean. This method removes x percent of the highest and lowest measurements in your series, creating a more stable average. We won't go any deeper into the trimmed mean for baseline use, but if you want to learn more about it I suggest you read Bob Newstadt's blog post over at <http://www.sqlteam.com/article/computing-the-trimmed-mean-in-sql>. Even though the article is pretty dated, it shows a method of calculating the trimmed mean using T-SQL.

Baseline Pitfalls

Hopefully the previous section convinced you that baselines are important, but before you go and capture every performance metric and convert it into a baseline measurement there are some pitfalls you will want to avoid.

Too Much Information

Even though you are free to baseline everything in your system, this is generally considered a bad idea. Gathering too much information can blind you in your search for the answers. If you have to compare 100,000 different metrics against your baseline every time a performance problem occurs, you are wasting time. The advice here is to keep your baselines small, including only performance metrics that matter the most for your system. For instance, you can include performance metrics related to Availability Groups, but if your system doesn't use this feature then there is no use in your including them.

Know Your Metrics

Another important aspect in the selection of performance metrics is understanding. If you do not understand what a performance metric represents it can be very difficult to formulate a correct conclusion, or it can even lead you in the wrong direction.

Focus on the Big Measurement Changes

When comparing measurements against a baseline, always focus on the big increases or decreases. Especially for Wait Statistics, very small increases in Wait Time (1%–2%) aren't a cause for concern. If one of your Wait Time measurements goes up 20%, that would be a good signal to start investigating.

Use Fixed Intervals

When capturing Wait Statistics information you should always use a fixed interval. If we were to capture Wait Times at random it would be almost impossible to build a reliable baseline. It would be like comparing apples against oranges. The best way to automate the capture of Wait Statistics information is by using the SQL Server Agent and setting it to a fixed interval, like every 15 minutes.

Building a Baseline for Wait Statistics Analysis

Now that we have familiarized ourselves with baselines, let's get to work and create a baseline we can use in our Wait Statistics analysis. As I mentioned at the beginning of this chapter, baselines are incredibly important if you want to analyze performance problems using Wait Statistics. Nobody has the same Wait Types and Wait Times compared to your system, so it's up to you to create a baseline you can compare against.

In this section I will show you a method I use to create, maintain, and compare baselines and measurements. This does not necessarily mean this is the right way to do it, and you might find other methods better suited to your needs.

Since we are going to capture SQL Server Wait Statistics measurements I prefer to store my measurements inside a separate database named "Baseline." This way my measurement information doesn't get stored somewhere between user tables. Since Wait

Statistics are logged at the SQL Server instance level, it makes sense to create a separate measurement database inside every SQL Server instance. Figure 3-4 shows you my baseline database inside the SQL Server Management Studio.




Figure 3-4. Baseline database

You can create the database yourself using the script in Listing 3-1, making sure to change the file locations. The database data file will be 1.5 GB when it gets created, which gives enough free space to capture weeks of Wait Statistics information.

Listing 3-1.

```
-- Create Baseline database
CREATE DATABASE [Baseline]
ON PRIMARY
(
    NAME = 'Baseline', FILENAME = 'E:\Data\baseline_data.mdf', SIZE =
    1536000KB, FILEGROWTH = 10%
)
LOG ON
(
    NAME = 'Baseline_log', FILENAME = 'E:\Log\baseline_log.ldf', SIZE =
    102400KB , FILEGROWTH = 10%
)
GO

ALTER DATABASE [Baseline] SET RECOVERY SIMPLE
GO
```

We will be using the `sys.dm_os_wait_stats` DMV as the source of our measurements, which means that the table that will hold our measurements must be able to handle the information returned from the DMV. We will not only store the Wait Types and Wait Times, but will also add additional information to enrich the data so that we can easily create multiple baselines.

Listing 3-2 shows the query you can use to create a table, named `WaitStats`, to hold the Wait Statistics information we will use for creating our baselines.

Listing 3-2. Create a Wait Statistics table

```
USE [BaseLine]
GO

CREATE TABLE WaitStats
(
  ws_ID INT IDENTITY(1,1) PRIMARY KEY,
  ws_DateTime DATETIME,
  ws_Day INT,
  ws_Month INT,
  ws_Year INT,
  ws_Hour INT,
  ws_Minute INT,
  ws_DayOfWeek VARCHAR(15),
  ws_WaitType VARCHAR(50),
  ws_WaitTime INT,
  ws_WaitingTasks INT,
  ws_SignalWaitTime INT
)
```

As you can read in the listing, we capture the Wait Type, Wait Time, Signal Wait Time, and the number of waiting tasks. We also capture the date and time when we log the Wait Statistics information. We also split the date and time into additional columns to segment the data, making it easier to build specific baselines based on a specific day, hour, month, and so forth, without having to convert the `datetime` data type every time.

Now that we have our table ready it is time to capture some Wait Statistics and insert them into our `WaitStats` table. Because the `sys.dm_os_wait_stats` DMV returns cumulative Wait Times, we have to use a method to only capture the difference in Wait Time between two capture moments. If we were to only capture the information directly from the `sys.dm_os_wait_stats` DMV we would always receive ever-increasing Wait Times, and that would make comparisons useless. There are two paths we can take to capture the change in Wait Time between two measurements, and both have their advantages and disadvantages.

The first method, which I call the reset method, will capture the Wait Statistics information from the `sys.dm_os_wait_stats` DMV and then reset the DMV using the `DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR)` command. The main advantage of this method is that it is very simple to use, as we only need to capture the information, reset it again, and start the same procedure at the next measurement. There is no need to calculate deltas, because after our first measurement the counters are reset to 0.

The disadvantage is that the DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR) command resets the information inside the sys.dm_os_wait_stats DMV. This means that you will lose the cumulative information inside the DMV, information you might not want to lose. Figure 3-5 illustrates this method of Wait Statistics capturing.



Figure 3-5. Capturing Wait Statistics using the reset method

The second option, which I named the delta method, involves not using the DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR) command, but rather calculating the difference, or delta, in Wait Time between two measurements. The advantage of not using the DBCC command is that you will not lose the cumulative Wait Times inside the sys.dm_os_wait_stats DMV. Its main disadvantage is that it is a lot more complex to calculate the deltas compared to the first method. It usually also involves a WAITFOR DELAY command inside the T-SQL script to set the interval. This might mean that, if you plan to capture Wait Statistics information using the SQL Server Agent, you could end up with a SQL Server Agent job that is running almost continuously. Figure 3-6 illustrates the delta option of capture Wait Statistics.

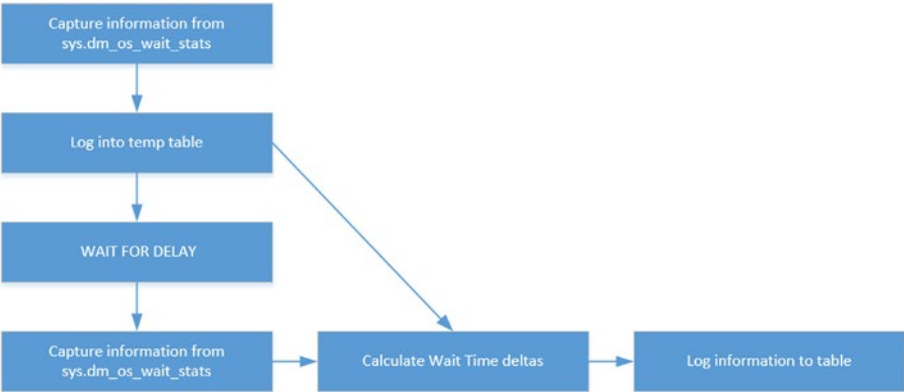


Figure 3-6. Capturing Wait Statistics using the delta method

There are more methods available with which to capture Wait Statistics information, but I most frequently see these two, or variations of them. What method you want to use is up to you, as in the end both will return the same results.

Reset Capture Method

The reset Wait Statistics capture method consists of a single T-SQL script that will capture the information from the `sys.dm_os_wait_stats` DMV followed by a reset of the counters inside the DMV. Listing 3-3 shows the T-SQL script you can use to fill the `WaitStats` table we created earlier.

Listing 3-3. Reset capture method

```
USE [Baseline]
GO

-- Insert Wait Stats into Baseline table
INSERT INTO WaitStats
SELECT
    GETDATE() AS 'DateTime',
    DATEPART(DAY,GETDATE()) AS 'Day',
    DATEPART(MONTH,GETDATE()) AS 'Month',
    DATEPART(YEAR,GETDATE()) AS 'Year',
    DATEPART(HOUR, GETDATE()) AS 'Hour',
    DATEPART(MINUTE, GETDATE()) AS 'Minute',
    DATENAME(DW, GETDATE()) AS 'DayOfWeek',
    wait_type AS 'WaitType',
    wait_time_ms AS 'WaitTime',
    waiting_tasks_count AS 'WaitingTasks',
    signal_wait_time_ms AS 'SignalWaitTime'
FROM sys.dm_os_wait_stats;

-- Clear sys.dm_os_wait_stats
DBCC SQLPERF ('sys.dm_os_wait_stats',CLEAR)
GO
```

Delta Capture Method

The delta capture method also consists of a single T-SQL script, but it is a little more complex than the reset capture method. It uses a temporary table to store the first measurement, then waits for 15 minutes, performs a second measurement, and calculates the deltas. The result is inserted into the `WaitStats` table. Listing 3-4 shows the T-SQL script you can use if you plan to use this method of collecting Wait Statistics metrics.

Listing 3-4. Delta capture method

```
USE [Baseline]
GO

-- Check if the temp table already exists
-- if it does drop it.
```

```

IF EXISTS
(
  SELECT *
  FROM tempdb.dbo.sysobjects
  WHERE ID = OBJECT_ID(N'tempdb..#ws_Capture')
)
DROP TABLE #ws_Capture;

-- Create temp table to hold our first measurement
CREATE TABLE #ws_Capture
(
  wst_WaitType VARCHAR(50),
  wst_WaitTime INT,
  wst_WaitingTasks INT,
  wst_SignalWaitTime INT
);

-- Insert our first measurement into the temp table
INSERT INTO #ws_Capture
SELECT
  wait_type,
  wait_time_ms,
  waiting_tasks_count,
  signal_wait_time_ms
FROM sys.dm_os_wait_stats;

-- Wait for the next measurement
-- In this case we will wait 15 minutes
WAITFOR DELAY '00:15:00'

-- Combine the first measurement with a new
-- measurement and calculate the deltas
-- Write the results into the WaitStats table
INSERT INTO WaitStats
SELECT
  GETDATE() AS 'DateTime',
  DATEPART(DAY,GETDATE()) AS 'Day',
  DATEPART(MONTH,GETDATE()) AS 'Month',
  DATEPART(YEAR,GETDATE()) AS 'Year',
  DATEPART(HOUR, GETDATE()) AS 'Hour',
  DATEPART(MINUTE, GETDATE()) AS 'Minute',
  DATENAME(DW, GETDATE()) AS 'DayOfWeek',
  dm.wait_type AS 'WaitType',
  dm.wait_time_ms - ws.wst_WaitTime AS 'WaitTime',
  dm.waiting_tasks_count - ws.wst_WaitingTasks AS 'WaitingTasks',
  dm.signal_wait_time_ms - ws.wst_SignalWaitTime AS 'SignalWaitTime'

```

```
FROM sys.dm_os_wait_stats dm
INNER JOIN #ws_Capture ws
ON dm.wait_type = ws.wst_WaitType;

-- Clean up the temp table
DROP TABLE #ws_Capture;
```

Using SQL Server Agent to Schedule Measurements

After selecting a capture method, we need to run the capture T-SQL script to fill our `WaitStats` table with Wait Statistics information. As described in the baseline pitfalls section earlier, it is very important to always perform your measurements at a fixed interval. This makes comparing measurements a lot easier, since you are always comparing the same time segments. The best way to do this is by using a SQL Server Agent job set to a fixed interval. The interval can be set to your choosing—the larger you set the interval, the smaller the number of time segments you can compare against. Setting the interval to be shorter will give you more time segments, but will also mean an increase in data that you need to store. I personally prefer to set my interval to 15 minutes. This gives me enough time segments to compare in most cases.

I won't go into details here about how you can create a SQL Server Agent job to capture Wait Statistics information, but I do want to point out how my job looks as an example you can use. I usually end up with a SQL Server Agent job with just one T-SQL script step. In this step I copy the capture script, depending on which method I want to use. Figure 3-7 shows a screenshot of my SQL Server Agent job.

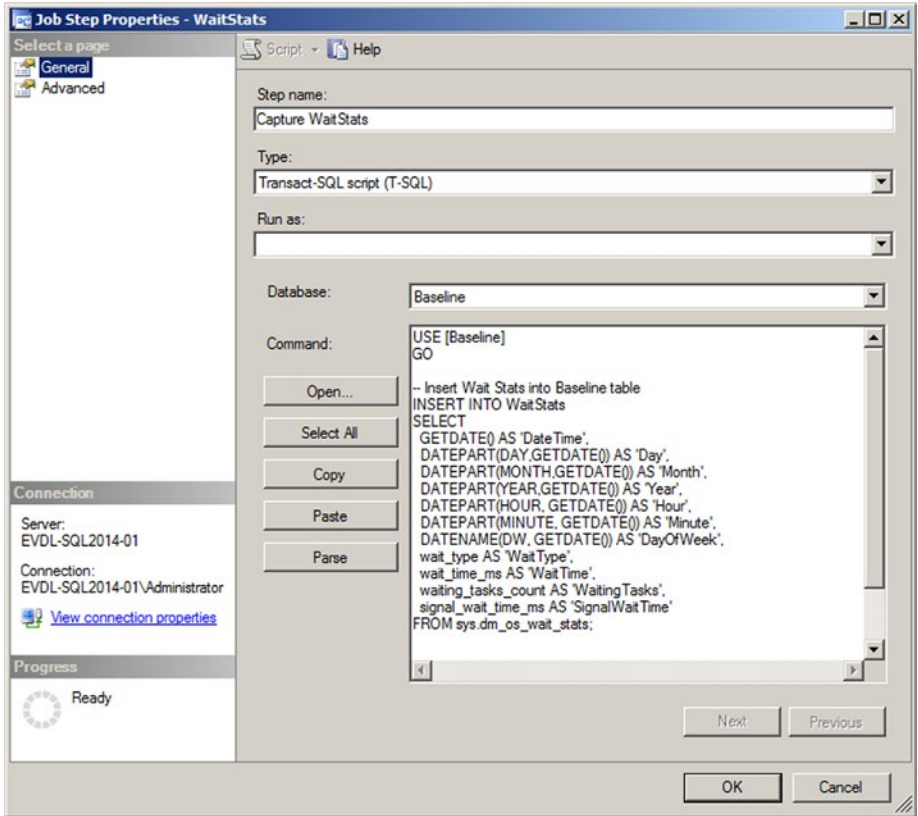


Figure 3-7. Capturing WaitStats SQL Server Agent job step

In this case I used the reset capture method to capture the Wait Statistics in my WaitStats table.

Figure 3-8 shows the schedule I use to capture the Wait Statistics on a fixed interval. As you can see, I have set it to every 15 minutes, every day.

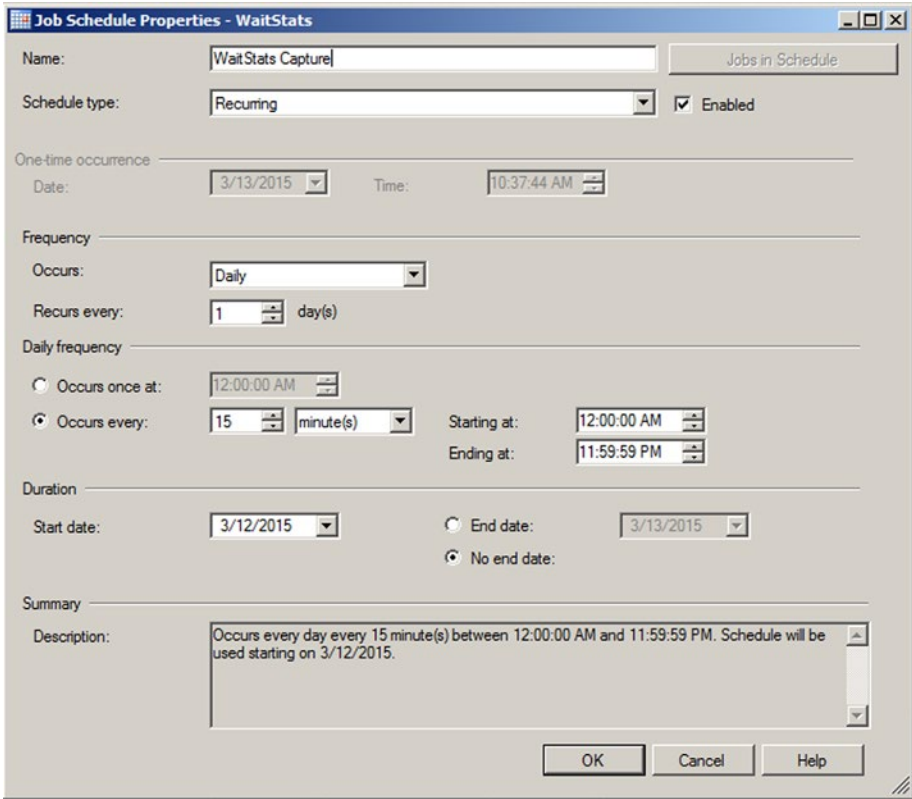


Figure 3-8. SQL Server Agent job schedule

Again, you are free to choose your own capture interval, but make sure to always capture at the same interval length.

After we have created a SQL Server Agent job to gather Wait Statistics information, we need to let it run for a while. The longer the job runs, the more information we gather, improving the quality of our baselines.

Wait Statistics Baseline Analysis

After letting the SQL Server Agent job collect Wait Statistics metrics for a while, we are ready to actually create some baselines. The way we do this is by querying the `WaitStats` table we created earlier. I will give you some examples of queries that will create a baseline you can compare against; these are not the only queries you can run, however, and I encourage you to experiment with different queries to return the information you are most interested in.

Before we get started with the baseline building I want to return to Figure 2-13 in Chapter 2. In this flowchart I showed you steps you can take to analyze resource waits that occur right now. Since we now have access to a baseline, we can add an extra step to the flowchart. Figure 3-9 shows how to complete the flowchart, including the baseline comparison step.

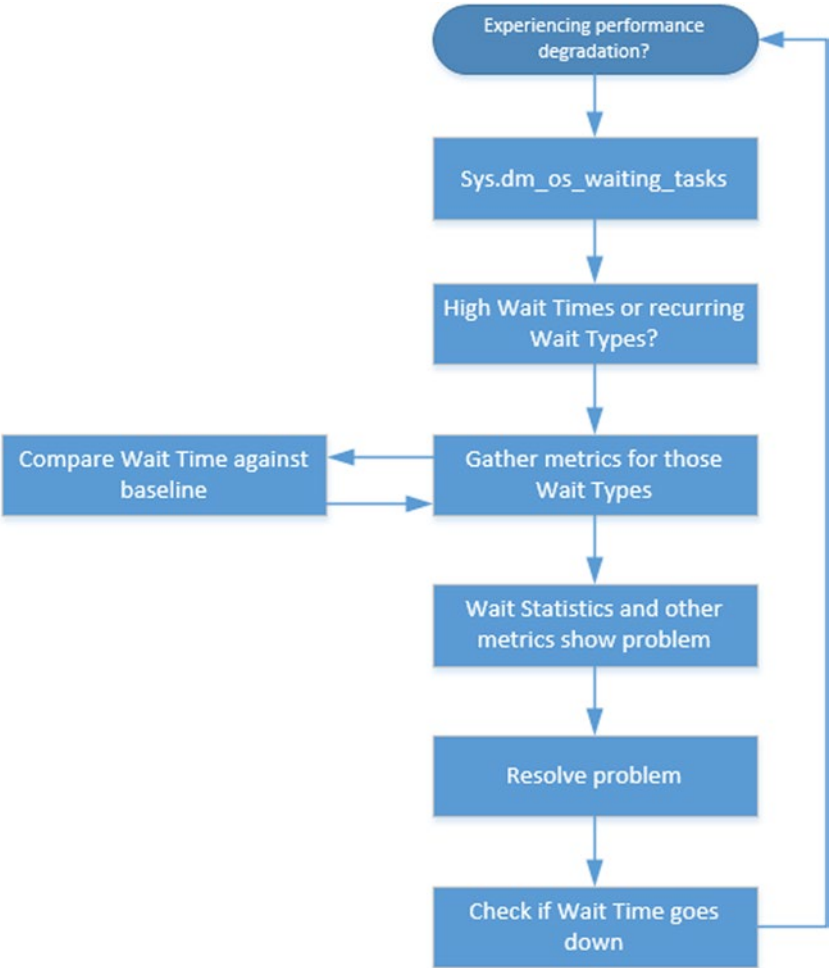


Figure 3-9. Complete Wait Statistics performance-analysis flowchart

The baselines you create are an extra input to the metrics you gather when looking at a performance problem. They are a very valuable input because they will show you information about the time the problem didn't exist.

Let's go through an example, using DBA Jim again, where we review all the steps of the flowchart shown in Figure 3-9. In this example I will show you queries that you can use against the `WaitStats` table so as to build a baseline that is useful for the performance-analysis process.

Tuesday, around 9 AM, DBA Jim receives a phone call that the daily reporting against the sales database is a lot slower than normal. The problem started around 8 AM, and users are still experiencing performance problems. The reports are part of a scheduled job that runs every workday, starting at 8 AM.

The first thing Jim does is query the `sys.dm_os_waiting_tasks` DMV using the following query:

```
SELECT * FROM sys.dm_os_waiting_tasks
ORDER BY session_id ASC;
```

Jim focuses on user sessions (normally higher than ID 50) but doesn't see any long Wait Times on any of the user sessions, as shown in Figure 3-10.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address
14	0x00000000E412748	18	0	4799389616	ONDEMAND_TASK_QUEUE	0x00000000388E278
15	0x0000000006012748	19	0	49452	CHECKPOINT_QUEUE	0x00000000038A9800
16	0x000000000C41C8C8	20	0	199506583	BROKER_EVENTHANDLER	NULL
17	0x00000000E430748	21	0	4799387245	BROKER_TRANSMITTER	NULL
18	0x000000001041C508	23	0	4799387245	BROKER_TRANSMITTER	NULL
19	0x000000001043A988	2060	0	4	PAGEIOLATCH_SH	0x00000010B9FDEF...
20	0x0000000006038DC8	4002	0	0	PAGEIOLATCH_EX	0x00000009CDFE00...
21	0x00000000104134C8	4519	0	6	PAGEIOLATCH_SH	0x00000010B9FD32...

Figure 3-10. `sys.dm_os_waiting_tasks`

After executing the query against the `sys.dm_os_waiting_tasks` DMV multiple times, Jim notices that the Wait Type `PAGEIOLATCH_SH` is returned every time he queries the DMV. Each time, the Wait Type is returned with a different session ID but with relatively low Wait Times.

Jim uses the same T-SQL script to capture Wait Statistics metrics into the `WaitStats` table as we discussed earlier in this chapter. Because Jim has access to historic Wait Statistics information, he decides to create a baseline of the `PAGEIOLATCH_SH` Wait Times. The first thing he does is view the `PAGEIOLATCH_SH` Wait Times of today, filtered to show measurements captured between 8 and 9 in the morning, using the query shown in Listing 3-5.

Listing 3-5. Show Wait Times for `PAGEIOLATCH_SH` between 8 and 9 AM today

```
-- PAGEIOLATCH_SH waits, today between 8 and 9 AM
SELECT
    CONVERT(VARCHAR(5), ws_DateTime, 108) AS 'Time',
    ws_WaitTime AS 'Wait Time'
FROM WaitStats
WHERE ws_WaitType = 'PAGEIOLATCH_SH'
AND (ws_Hour >= 8 AND ws_Hour < 9)
AND CONVERT(VARCHAR(5), ws_DateTime, 105) = CONVERT(VARCHAR(5), GETDATE(), 105)
```

The query returned the results shown in Figure 3-11.

	Time	Wait Time
1	08:00	1528749
2	08:15	1828749
3	08:30	1658974
4	08:45	1698547

Figure 3-11. PAGEIOLATCH_SH Wait Times of today

Now that Jim has the Wait Times for today of the PAGEIOLATCH_SH Wait Type, the next step is to create a baseline from the historic measurements of the PAGEIOLATCH_SH Wait Type so he can compare today's measurements against the baseline. Jim uses the query shown in Listing 3-6 to build his baseline.

Listing 3-6. PAGEIOLATCH_SH baseline

```
-- Baseline between 8 and 9 on workdays
-- Not including measurements done today
SELECT
    CONVERT(VARCHAR(5), ws_DateTime, 108) AS 'Time',
    AVG(ws_WaitTime) AS 'Baseline'
FROM WaitStats
WHERE ws_WaitType = 'PAGEIOLATCH_SH'
AND ws_DayOfWeek IN ('Monday', 'Tuesday', 'Wednesday', 'Thursday','Friday')
AND (ws_Hour >= 8 AND ws_Hour < 9)
AND CONVERT(VARCHAR(5), ws_DateTime, 105) < CONVERT(VARCHAR(5), GETDATE(),
105)
GROUP BY CONVERT(VARCHAR(5), ws_DateTime, 108);
```

This query builds a baseline with the following characteristics: return the average Wait Time of PAGEIOLATCH_SH Wait Types captured on a workday between 8 AM and 9 AM, excluding today. The reason to exclude today is that the measurements that were performed today, during the performance problem, might impact the average. Another suggestion could be to filter only data captured in the last x weeks so as to limit the amount of data that needs to be calculated in the average.

The results of the query shown in Listing 3-6 can be seen in Figure 3-12.

	Time	Baseline
1	08:00	313038
2	08:15	391444
3	08:30	498923
4	08:45	570782

Figure 3-12. PAGEIOLATCH_SH baseline

As you can immediately see when you compare the Wait Times in Figure 3-11 and Figure 3-12, the measurements done today are a lot higher than those in the historic baseline. To make it a little easier to see the difference, I created a graph of both measurements, as shown in Figure 3-13.




Figure 3-13. Baseline comparison graph for the PAGEIOLATCH_SH Wait Type

Because there is such a difference in Wait Times for the PAGEIOLATCH_SH Wait Type between the baseline and today, Jim believes the PAGEIOLATCH_SH Wait Type needs further investigation.

We will take a detailed look at the PAGEIOLATCH_SH Wait Type in Chapter 8, “Latch Wait Types,” but to give you a (very) short explanation, long PAGEIOLATCH_SH waits can indicate storage problems.

To investigate further, Jim starts the Windows Performance Monitor to look at metrics related to the storage subsystem, and in particular the disk latency counters. As you can see in Figure 3-14, the latency on the disk where the database data file resides peaks to very high values, more than 4,000 milliseconds! For SQL Server to perform optimally, the disk latency should be as low as possible, and at least below 20 milliseconds.




Figure 3-14. Disk-read latency

With both the Wait Statistics baseline information and the Perfmon metrics, Jim believes the problem is storage related and contacts the storage administrator. The metrics Jim collected also help the storage administrator, since he can compare his storage-related measurements against those Jim supplied. The storage administrator confirms there is a problem related to the disk that contains the sales database, and solves the problem by replacing a faulty disk in the disk array. After the disk gets replaced the disk latency returns to a 6 millisecond average, and the high latency peaks disappear. Jim queries the Wait Times again from the `WaitStats` table after the disk is replaced and notices the Wait Times for the `PAGEIOLATCH_SH` Wait Type are close to the baseline values again. Users also inform Jim that the reports are running normally again.

During this example Jim went through all the steps of the Wait Statistics performance-analysis flowchart shown in Figure 3-9:

1. Users experience performance degradation while running reports.
2. Jim queries the `sys.dm_os_waiting_tasks` DMV to find out if there are high Wait Times or frequently recurring Wait Types. The `PAGEIOLATCH_SH` Wait Type seems to be recurring frequently.
3. Jim gathers metrics by capturing the `PAGEIOLATCH_SH` Wait Times of today and comparing them to the baseline. He also gathers additional metrics from Perfmon.
4. All the metrics show Jim that the problem is most likely storage related, and Jim contacts the storage administrator.

5. The storage administrator replaces a broken disk in the array. Storage latency values drop to 6 milliseconds.
6. Jim checks the Wait Times of the PAGEIOLATCH_SH Wait Type again and confirms that they are close to the baseline values.

Even though this example might seem very simple, it is actually based on a performance problem I encountered in the real world. Using the steps from the Wait Statistics performance-analysis flowchart combined with the baseline metrics, I was able to identify and solve the problem very quickly.

In the example I showed you the query in Listing 3-6 that creates a baseline for the PAGEIOLATCH_SH Wait Type. This query is just an example of what you can use against the WaitStats table. You can modify it to suit your own needs; for instance, you can choose to not limit the results for weekdays, and only show average Wait Times captured on a specific day. Or you could request the actual Wait Times on a specific date.

If you are capturing Wait Statistics measurements for a long period of time, it might be a good idea to split the results into multiple tables for easier and faster querying. For instance, you could use the following query to insert all the Wait Statistics measurements done in March into their own table:

```
SELECT *
INTO WaitStats_March
FROM WaitStats
WHERE ws_Month = 3;
```

This also give us options to compare specific Wait Times during different periods of time by joining the different tables together. Figure 3-15 shows the tables of my baseline database that I usually end up with, sorting the data per month.

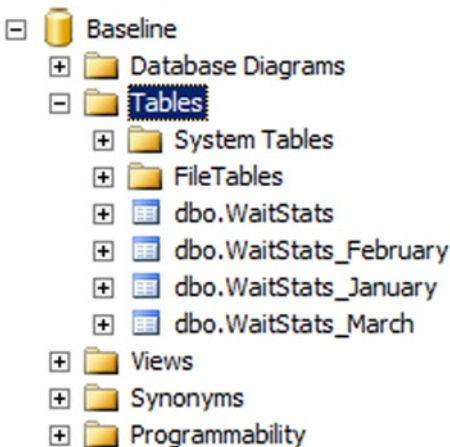


Figure 3-15. Wait Statistics measurement split per month

You can decide how you want to split the measurements yourself; maybe you want to store the Wait Statistics measurements in a separate table for every application version you use, or store all the measurements of a specific Wait Type in a separate table. The choice is yours.

This chapter hopefully gave you some ideas on how to store Wait Statistics measurements and create baselines from those measurements. I tried to avoid telling you exactly what to do and how to do it, because I believe one single approach doesn't work for everybody. You will need to write and adjust your own queries to create the baselines you are interested in, but I hope this chapter showed you the foundations to further build upon.

Summary

In this chapter we took a close look at baselines from both a theoretical and a practical point of view. Baselines are incredibly important for any type of performance analysis you perform. In the case of Wait Statistics, baselines are frequently required if you want to troubleshoot SQL Server-related performance problems. Since Wait Statistics are unique for your system, there is only one method by which to compare Wait Times—baselines.

I gave you some examples and T-SQL scripts to create your own Wait Statistics baseline table so you can start capturing Wait Statistics information right now. We also went through an example of how you can query that baseline information and compare it to actual measurements to troubleshoot a performance-related incident.

PART II



Wait Types

Introduction

In the first part of this book we focused on the foundations of SQL Server Wait Statistics. I believe that understanding these foundations are necessary before you move on to reading about specific Wait Types, knowing how Wait Types and Wait Times are generated will help you resolve, or lower, problematic waits faster.

In this second part we won't go back into detail about how Wait Statistics are generated, or explain the various DMV's you can use to analyze and troubleshoot them. Instead we will focus on specific Wait Types and give you the tools needed to troubleshoot them.

Because there are over 700 different Wait Types in SQL Server 2014 it is impossible to include them all in this book. To keep this book readable I included only the most common or most problematic Wait Types and categorized them to the areas they are related to. Because many Wait Types potentially touch upon different categories, you might find certain Wait Types in a different category than you expected. This also impacts the suggestions on how you can lower, or resolve, Wait Times of the specific Wait Types. In a lot of cases there are many different ways to lower specific Wait Times. Because of this I included the most common methods, or those with the highest chance of success, for resolving those Wait Times.

I have tried to include examples for generating that specific Wait Type for all Wait Types included in this book. Some Wait Types though, will not include an example on how you could generate them in a test environment.

In those cases the examples were either too random or too complex to include. The examples were created in my test SQL Server Instance and should work on your test environment as well, but because Wait Statistics are heavily influenced by your hardware and SQL Server configuration, your results might be different compared to my examples.

I want to make one final, important, note about the examples included to generate specific Wait Types. Hopefully it goes without saying but, do not perform these examples on a production environment!

CHAPTER 4



CPU-Related Wait Types

Processors have evolved enormously in the last few decades, and processor manufacturers, like Intel or AMD, manage to build faster processors on a yearly basis. And while the speed of processors is hitting a ceiling, the number of cores manufacturers manage to build inside their processors has only grown. At the time of writing this book, you can buy a single processor with 18 cores inside to power your system. Processors are also one of the more difficult parts of your system to replace. While you can expand your system's memory relatively simply, replacing a processor for one that is faster or has more cores frequently requires you to change your system's motherboard as well due to CPU socket incompatibility. This means we are usually stuck with our processors until we replace the system altogether.

Processors are also very important for SQL Server. Higher processor speeds will accelerate processor-related instructions, and more cores means more schedulers that SQL Server can use to execute requests. But even all these upgrades in speed and cores cannot prevent the fact that we sometimes have to wait on processor resources. In this chapter we will take a look at some of those Wait Types that have a relation with your system's processor.

CXPACKET

The first CPU-related Wait Type is also the most common Wait Type in SQL Server instances that run with the default, out-of-the-box, SQL Server configuration. It is also one of the most misunderstood Wait Types and sometimes doesn't even need lowering in order to make your queries perform faster; as a matter of fact, lowering CXPACKET Wait Times can sometimes degrade the performance of your queries!

What Is the CXPACKET Wait Type?

The CXPACKET Wait Type occurs whenever a query is being executed in parallel instead of serially. Parallel queries can have a performance advantage compared to serial queries, if the work can be divided among multiple worker threads. The advantage is bigger for

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-1-4842-1139-7_4](https://doi.org/10.1007/978-1-4842-1139-7_4)) contains supplementary material, which is available to authorized users.

queries that are returning large result sets; queries that return only a few rows benefit far less from parallelism, and in many cases parallelism can slow down those queries. This doesn't mean we should turn off parallelism immediately, as I have yet to see a true OLTP database where every query only returns a handful of rows. Many systems have to deal with a mixed workload, usually dealing with many short queries but also large, longer-running, reporting queries.

Parallel queries will use multiple worker threads to execute a request. Along with the worker threads that are created to perform the work requested, a parallel query will also use a 0 thread. This 0 thread's task is to coordinate the work of the other worker threads. While the 0 thread is waiting for the other worker threads to finish the work they were assigned to perform, it will record Wait Times of the CXPACKET Wait Type. To make understand this relation a little bit better, take a look at Figure 4-1.



Figure 4-1. Parallel query threading

As soon as the SQL Server Query Optimizer decides on an execution plan that uses parallelism, you will see CXPACKET waits occur. This can be completely normal and is nothing to worry about if you are expecting your queries to run in parallel and they are performing as expected. In those cases you can ignore long Wait Times on the CXPACKET Wait Type. There are, however, cases where you don't want to use parallelism, or when parallelism is negatively impacting the performance of your queries because of skewed workloads.

Because the CXPACKET Wait Type is directly related to the parallelism settings of your SQL Server instance, we can influence it relatively easily by adjusting these settings. We can find the parallelism settings in the **Server Properties - Advanced** section of your SQL Server instance, as shown in Figure 4-2.

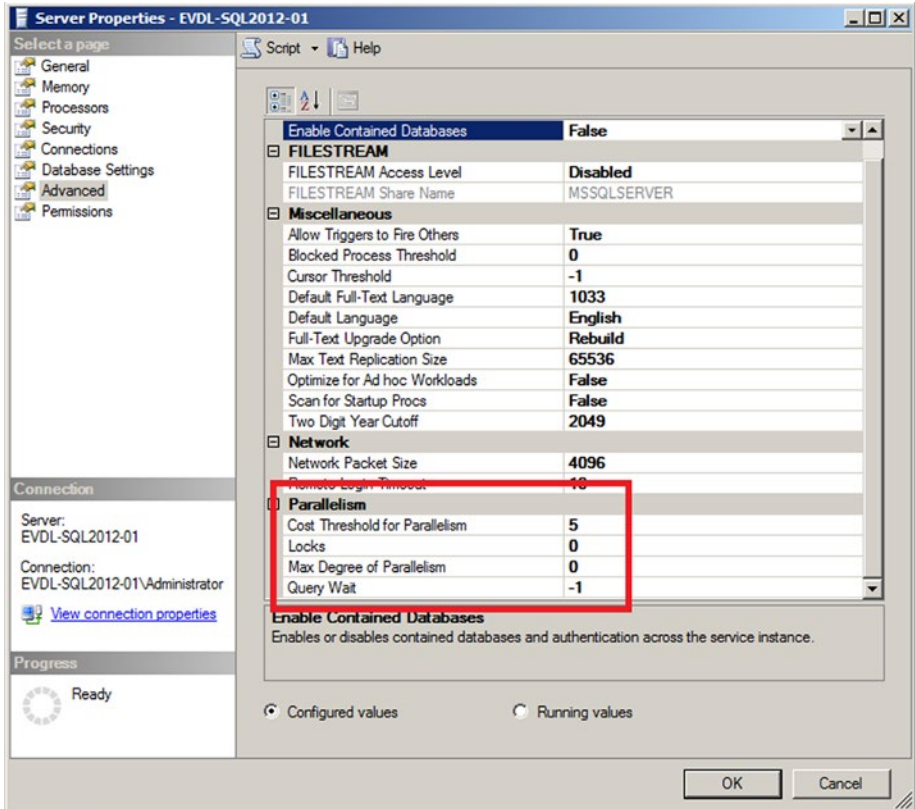


Figure 4-2. Parallelism configuration

Of these settings, the **Cost Threshold for Parallelism** and **Max Degree of Parallelism** settings impact parallel queries the most.

The **Cost Threshold for Parallelism** setting configures the cost threshold of when a query will be considered to be run in parallel by the Query Optimizer. If a serial query has a cost higher than the value configured in the **Cost Threshold of Parallelism**, the Query Optimizer might decide to generate a parallel plan instead of a serial one. By default, the setting has a value of 5 and can be configured to have a value between 0 and 32767.

The **Max Degree of Parallelism** setting configures the number of schedulers used when executing a parallel plan. By default, this setting is configured to be 0, which means all available schedulers can be used when a parallel plan is executed.

Lowering CXPACKET Wait Time by Tuning the Parallelism Configuration

There are various methods you can use to lower CXPACKET Wait Times, but before you go and use them you have to be sure that CXPACKET waits are actually causing you problems. Like I said earlier, CXPACKET waits are completely normal whenever you have parallelism enabled for your SQL Server instance. One solution I read frequently on the Internet forums is to disable parallelism by setting the **Max Degree of Parallelism** option to a value of 1. In most cases this is not a good idea. Disabling parallelism will make the CXPACKET waits go away completely, but some of your queries might be performing a lot worse since they cannot be run in parallel anymore.

A better approach to lowering CXPACKET waits is to tune the **Cost Threshold for Parallelism** and **Max Degree of Parallelism** options so they match with your workload. This way you can make sure only the queries that benefit the most from parallelism will be run in parallel. A way to find this parallelism sweet spot is by comparing the runtime of a query when it ran serially versus in parallel. You should generally focus on queries that access a lot of information and have a longer runtime in general, as those will be the queries that benefit the most from parallelism.

Consider this example where we have a query against the AdventureWorks database that requests information from the Sales.SalesOrderDetail table:

```
SELECT *
FROM Sales.SalesOrderDetail
ORDER BY CarrierTrackingNumber DESC;
```

We can check if this query would be a candidate to be run in parallel by checking the estimated cost of the query. To view this information we need to take a look at the estimated execution plan for if the query were to be run serially. To make sure the query runs serially we must add the query option `MAXDOP 1`. We are also interested in the runtime of the query, so we add the `SET STATISTICS TIME ON` option to the query:

```
SET STATISTICS TIME ON

SELECT *
FROM Sales.SalesOrderDetail
ORDER BY CarrierTrackingNumber DESC
OPTION (MAXDOP 1);

SET STATISTICS TIME OFF
```

Figure 4-3 shows the estimated cost of the query when run serially.

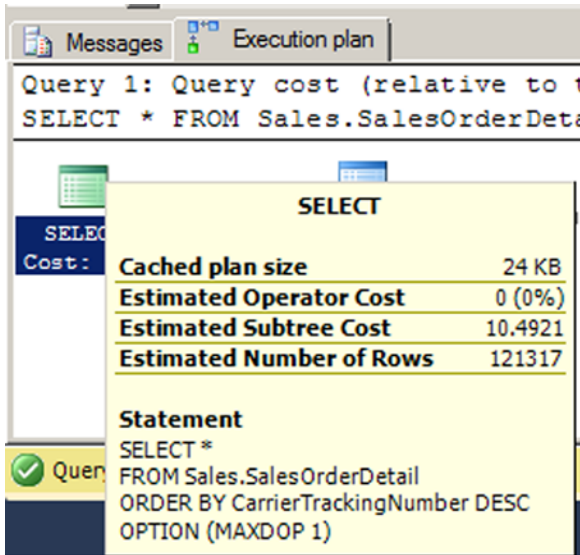


Figure 4-3. Estimated cost of the query with MAXDOP 1

In this case the estimated cost is 10.4921 on my test SQL Server. When I executed the query the execution time was 2256 milliseconds.

Because the Cost Threshold for Parallelism is still configured on the default value of 5 on my test server, I am pretty sure the query would be run in parallel if I were to remove the MAXDOP query hint.

Figure 4-4 shows the actual execution plan after running the query without the MAXDOP 1 option.




Figure 4-4. Actual execution plan without MAXDOP 1 option

As you can see, the query ran using parallelism, just as we expected, since the estimated cost was higher than the value configured in the Cost Threshold for Parallelism option. The execution time of the query with parallelism was 1959 milliseconds. If we take a look at the properties of the SELECT operation in the actual execution plan, we can view some additional information, as shown in Figure 4-5.

Query 1: Query cost (relative to
 SELECT * FROM Sales.SalesOrderDet

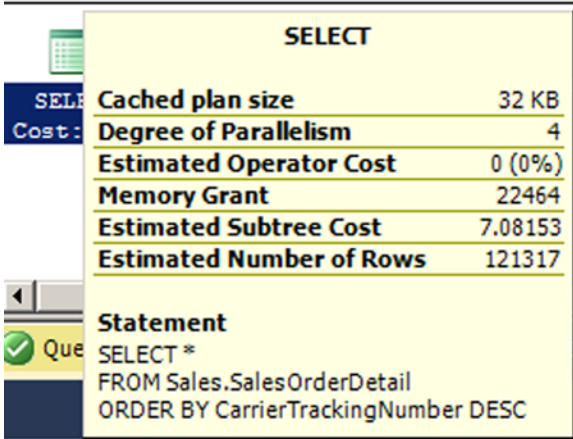


Figure 4-5. SELECT operation properties

The properties of the SELECT operation show us that the query was executed using four threads. The estimated cost went down to 7.08153.

Even though the estimated cost went down, the improvement in execution time is pretty small for this query. We could change the Cost Threshold for Parallelism value to a higher number than the default of 5. This way we are making sure relatively small queries like the one in this example don't use parallelism but that heavy reporting queries do.

Another setting to keep in mind is the Max Degree of Parallelism option. When it is set at its default of 0 all available schedulers can be used when a query runs in parallel. Using more schedulers doesn't necessarily mean the query executes faster though. The benefits of using more schedulers slowly get smaller after using more than 8. Microsoft recommends the following configuration in KB2806535:

- For servers with more than 8 cores, set the Max Degree of Parallelism option to 8.
- For servers with less than 8 cores, set the Max Degree of Parallelism option to 0 or to the amount of cores in your server.

This is a general recommendation, and your mileage may vary. The setting of both the Cost Threshold for Parallelism and Max Degree of Parallelism options highly depends on the workload of your system and requires careful testing to find out what works for you and what doesn't. They will impact your CXPACKET Wait Time though, so compare your CXPACKET Wait Times against a baseline after changing the Cost Threshold for Parallelism or Max Degree of Parallelism options to measure the impact of the change.

Lowering CXPACKET Wait Time by Resolving Skewed Workloads

A skewed workload means that all of the worker threads do not receive the same amount of work to perform. This is not an optimal situation, because if one worker thread has to do most of the work while another only a little bit, Thread 0 still has to wait for the longest-running worker thread to complete, logging CXPACKET waits as the time it is waiting. Figure 4-6 shows an abstract example of a skewed workload.




Figure 4-6. *Skewed parallel-query threading*

If we could give some of Thread 2's work to Thread 3, the query would probably perform faster, resulting in lower CXPACKET Wait Times.

We can view the thread distribution in the actual number of rows property of the parallel operation in the actual execution plan. Figure 4-7 shows the properties of a clustered index scan that has been performed using parallelism. The operation occurred in the example query we used in the previous section against the `Sales.SalesOrderDetail` table.

Clustered Index Scan (Clustered)	
Misc	
Actual Execution Mode	Row
Actual Number of Batches	0
Actual Number of Rows	121317
Thread 0	0
Thread 1	31355
Thread 2	27980
Thread 3	30385
Thread 4	31597
Actual Rebinds	0
Actual Rewinds	0

Figure 4-7. Parallel thread distribution

In this example we see that the clustered index scan returned 121.317 rows that were distributed among four threads (notice that Thread 0, the coordination thread, doesn't process any rows). The distribution of the number of rows is pretty even in this case, so we probably aren't running into a skewed workload problem.

Skewed workloads are frequently caused by outdated statistics. If the Query Optimizer believes there are fewer (or more rows) in the table than there actually are, it can distribute the work unevenly across the threads. Make sure to regularly perform maintenance on your statistics to prevent skewed workloads.

CXPACKET Summary

The CXPACKET Wait Type is directly related to the usage of parallelism during query execution. If you allow queries to be run using parallelism you will always see CXPACKET waits. Normally this is nothing to worry about, so avoid the knee-jerk reaction to turn off parallelism completely. Instead, focus on tuning the Max Degree of Parallelism and Cost Threshold for Parallelism options so that the thresholds are high enough that your large queries can benefit from using parallelism but your small queries do not experience a negative impact. Also, avoid skewed workloads by making sure your statistics are up-to-date.

SOS_SCHEDULER_YIELD

Just like CXPACKET, SOS_SCHEDULER_YIELD is a Wait Type that will frequently show up in the Top 10 of total Wait Time on your system. And just like the CXPACKET Wait Types, SOS_SCHEDULER_YIELD Wait Times do not necessarily indicate that there is a problem with your SQL Server instance. SOS_SCHEDULER_YIELD waits occur as soon as you start running queries on your SQL Server instance, and they are closely related to SQL Server scheduling.

What Is the SOS_SCHEDULER_YIELD Wait Type?

Before we can answer what the `SOS_SCHEDULER_YIELD` Wait Type means, we have to go back to Chapter 1 in this book, where we discussed SQL Server scheduling. Remember that the SQLOS uses its own cooperative non-preemptive scheduling model to make sure Microsoft Windows processes do not interrupt SQL Server's own processes? The `SOS_SCHEDULER_YIELD` Wait Type has a direct relation with the SQLOS's cooperative, non-preemptive scheduling model. To make it a little bit easier to understand, I have included Figure 4-8, which should be familiar to you as it represents a scheduler that we discussed in Chapter 1.




Figure 4-8. Scheduler and its phases and queues

If you remember from Chapter 1, worker threads move through the different phases and queues in a fixed order. Generally, a worker thread starts on the Waiter List while it waits for resources, it then moves to the Runnable Queue waiting for its turn to be run on the processor, and finally receives processor time to execute its request, receiving the “RUNNING” state. If the worker thread needs additional resources while it is in the “RUNNING” state, the worker thread moves back to the Waiter List and starts a new trip through the different queues and phases.

There is one exception to this behavior and it occurs when a worker thread is in the “RUNNING” state and doesn’t need additional resources to complete its work. If the SQLOS let a worker thread run on the processor for as long as it didn’t need any additional resources, the processor could be “hijacked” by one single worker thread for an infinite amount of time. To make sure a situation like this cannot occur, the scheduler gives every worker thread a specific slice of time in which they need to perform their work. We call this slice of time a quantum, and it is a fixed, unchangeable, 4 milliseconds. If a worker thread spends its quantum it has to yield the processor, and it then moves back to the bottom of the Runnable Queue. It will skip the Waiter List because the worker thread doesn’t need additional resources. While the worker thread is waiting to move back to the processor again, the `SOS_SCHEDULER_YIELD` Wait Type is recorded. Figure 4-9 shows this behavior.

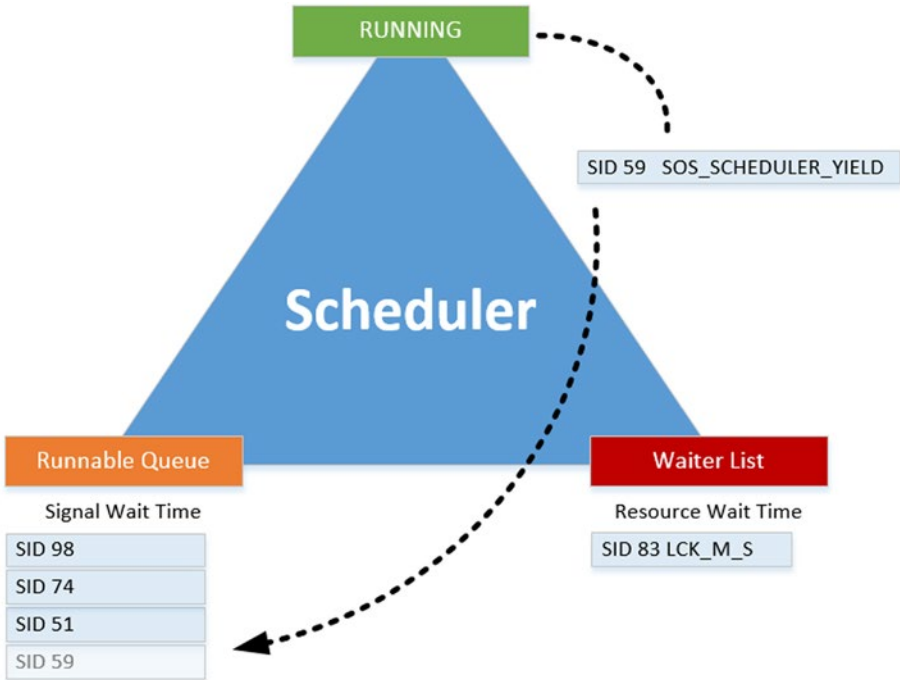


Figure 4-9. Worker thread voluntarily yielding the processor

As you can probably figure out, worker threads are voluntarily yielding all the time, especially on long-running queries where there is no need for additional resources. But keep in mind that Wait Times for the `SOS_SCHEDULER_YIELD` Wait Type will only be logged if the worker thread actually had to wait in the Runnable Queue. If there is no other worker thread in front of the yielding worker thread, it will move directly back to the processor without waiting (it will still move through the Runnable Queue though). To show you an example of this, I executed the queries below against the AdventureWorks database on my test SQL Server, where there is no concurrency whatsoever:

```
-- Clear Wait Stats
DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR);

-- Simple select
SELECT *
FROM Sales.SalesOrderDetail
ORDER BY CarrierTrackingNumber DESC;

-- Check for SOS_SCHEDULER_YIELD waits
SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'SOS_SCHEDULER_YIELD';
```

Figure 4-10 shows the results of this query against the `sys.dm_os_wait_stats` DMV.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	SOS_SCHEDULER_YIELD	60	0	0	0

Figure 4-10. `SOS_SCHEDULER_YIELD` waits

As you can see in Figure 4-10, the query against the AdventureWorks database encountered the `SOS_SCHEDULER_YIELD` Wait Type 60 times during execution. It didn't have to spend any time waiting for another worker thread in the Runnable Queue since this was the only query running at the time. If it had spent any time waiting for another worker thread, the `wait_time_ms` column would have returned a value higher than 0.

As I said at the start of this section, the `SOS_SCHEDULER_YIELD` Wait Type is generally not a cause for concern. If, however, the Wait Times are significantly higher than those in your baseline, it can be a reason to perform some additional research. There are basically three situations you can encounter when dealing with `SOS_SCHEDULER_YIELD` waits, as shown in Figure 4-11.

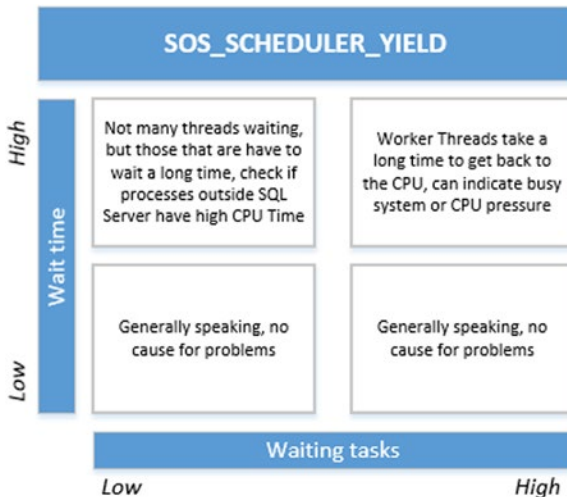


Figure 4-11. `SOS_SCHEDULER_YIELD` situations

Let's take a look at how we can analyze and resolve the SQL Server CPU pressure problem.

Lowering SOS_SCHEDULER_YIELD Waits

If you are experiencing higher than normal SOS_SCHEDULER_YIELD Wait Times and a large number of waits, you could, potentially, have a CPU-related problem on your system. To lower the SOS_SCHEDULER_YIELD waits, we are going to focus on the top right section of Figure 4-11, where there are a large amount of waiting tasks and high Wait Times.

If you are experiencing high Wait Times for the SOS_SCHEDULER_YIELD Wait Type together with a large amount of waiting tasks, you can assume you have a very busy SQL Server instance. Worker threads will yield, but it will take them a long time to get back on the processor again because there are many other threads waiting in the Runnable Queue. As we discussed earlier in Chapter 1, the Runnable Queue is a first in, first out list, meaning that the more worker threads that are waiting inside the Runnable Queue, the longer it takes for worker threads to move through it. You will usually see a high CPU usage on the system by the SQL Server process.

To show you an example of this problem, we will use the Ostress utility to execute a specific query simultaneously from a number of threads. The Ostress utility is part of the RML Utilities for SQL Server, which you can download here:

<https://support.microsoft.com/en-us/kb/944837>.

The first thing we are going to do is save the following query in a .sql file on the test server:

```
WHILE (1=1)
  BEGIN
    SELECT COUNT(*)
    FROM Sales.SalesOrderDetail
    WHERE SalesOrderID BETWEEN 45125 AND 54185
  END;
```

This query will count the number of rows between two SalesOrderIDs in the Sales.SalesOrderDetail table of the AdventureWorks database. It will do this in an endless loop.

After saving the query we start the Ostress utility using the following command:

```
"C:\Program Files\Microsoft Corporation\RMLUtils\ostress.exe" -E
-dAdventureWorks -i"C:\sos_scheduler_yield.sql" -n20 -r1 -q
```

This starts the Ostress utility, which connects to the AdventureWorks database and executes the sos_scheduler_yield.sql script using 20 threads.

As soon as we start Ostress, the CPU of the test SQL Server hits 100%, as shown in Figure 4-12.




Figure 4-12. Impact of Ostress on the CPU

As you can see in Figure 4-12, the CPU load is generated from the sqlservr#1 process, which happens to be the SQL Server Instance we are running the Ostress query against.

If we were to query the `sys.dm_os_waiting_tasks` DMV to check if the `SOS_SCHEDULER_YIELD` Wait Type is responsible for the CPU usage we would be in for a surprise, as you can see in Figure 4-13.

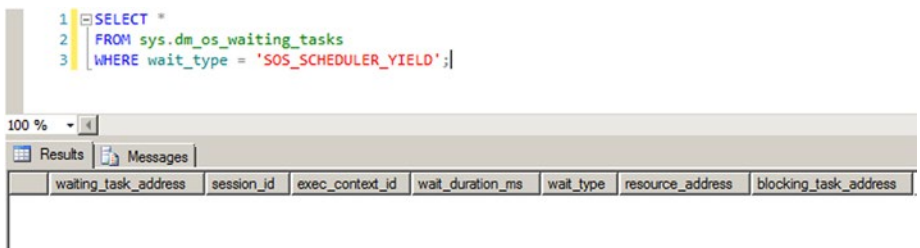


Figure 4-13. No `SOS_SCHEDULER_YIELD` waits occurring

This is the tricky bit of the `SOS_SCHEDULER_YIELD` Wait Type, as it frequently won't get returned by the `sys.dm_os_waiting_tasks` DMV—another reason to capture and use that Wait Statistics baseline!

To show that the high CPU usage is related to the `SOS_SCHEDULER_YIELD` Wait Type, we have to take a look at the cumulative Wait Statistics DMV, `sys.dm_os_wait_stats`. We can use the following query to show the top 10 Wait Types ordered by Wait Time while we run the `Ostress` utility (we can reset the DMV before starting the `Ostress` utility to keep the numbers small):

```
SELECT TOP 10 *
FROM sys.dm_os_wait_stats
ORDER by wait_time_ms DESC;
```

The results of this query are shown in Figure 4-14.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	SOS_SCHEDULER_YIELD	6667	107236	34	107219
2	XE_TIMER_EVENT	4	15004	5001	15004
3	REQUEST_FOR_DEADLOCK_SEARCH	3	15000	5000	15000
4	LAZYWRITER_SLEEP	14	14136	1027	124
5	DIRTY_PAGE_POLL	138	13997	104	2
6	LOGMGR_QUEUE	107	13900	132	0
7	HADR_FILESTREAM_IOMGR_IOCOMPLETION	27	13737	520	201
8	SQLTRACE_INCREMENTAL_FLUSH_SLEEP	3	12000	4000	0

Figure 4-14. Top 10 Wait Types during `Ostress` execution

As you can see, the number one Wait Type, by far, is `SOS_SCHEDULER_YIELD` with a pretty high amount of `waiting_tasks` and total `wait_time`.

If you were to experience this problem with a production SQL Server instance the first thing you should focus on are those very small, very quick queries like the ones we executed in this example. Has the volume of those queries increased? Has the number of user connections to the SQL Server executing those queries increased? Those are two quick questions you should ask and check. A sudden growth in transactions or user connections can lead to high `SOS_SCHEDULER_YIELD` Wait Times.

Another cause of high `SOS_SCHEDULER_YIELD` waits, together with very high CPU usage, can be a phenomena called *spinlock contention*. Spinlocks are defined by Microsoft as “lightweight synchronization primitives which are used to protect access to data structures” and are a very advanced topic. Appendix II, at the back of this book, goes into a little bit more detail about spinlocks for those who are interested in learning more about them.

Very large, very complex queries can also lead to higher `SOS_SCHEDULER_YIELD` Wait Times. Try looking for active queries that consume a lot of CPU time and have complex calculations or data-type conversions inside them. One query I use frequently to identify CPU-heavy queries is the one in Listing 4-1.

Listing 4-1. Detect expensive CPU queries

```
SELECT TOP 10
  QText.TEXT AS 'Query',
  QStats.execution_count AS 'Nr of Executions',
  QStats.total_worker_time/1000 AS 'Total CPU Time (ms)',
```

```

QStats.last_worker_time/1000 AS 'Last CPU Time (ms)',
QStats.last_execution_time AS 'Last Execution',
QPlan.query_plan AS 'Query Plan'
FROM sys.dm_exec_query_stats QStats
CROSS APPLY sys.dm_exec_sql_text(QStats.sql_handle) QText
CROSS APPLY sys.dm_exec_query_plan(QStats.plan_handle) QPlan
ORDER BY QStats.total_worker_time DESC;

```

The results of the query in Listing 4-1 on my test SQL Server can be seen in Figure 4-15.

	Query	Nr of Executions	Total CPU Time (ms)	Last CPU Time (ms)	Last Execution	Query Plan
1	WHILE (1=1) BEGIN SELECT COUNT(*) FRO...	132520	1503902	0	2015-03-23 21:47:10.107	<Show Plan>
2	SELECT * FROM Sales.SalesOrderDetail ORDER BY ...	26	60577	2242	2015-03-23 20:00:44.850	<Show Plan>
3	DBCC SQLPERF(sys.dm_os_wait_stats, CLEAR) SE...	16	35492	2168	2015-03-23 20:12:57.167	<Show Plan>
4	WHILE (1=1) BEGIN SELECT COUNT(*) FRO...	1155	12816	0	2015-03-23 20:54:30.110	<Show Plan>
5	- Simple select SELECT COUNT(*) FROM Sales Sale...	1002	12564	5	2015-03-23 20:52:07.490	<Show Plan>
6	WHILE (1=1) BEGIN SELECT COUNT(*) FRO...	509	6045	0	2015-03-23 20:54:30.107	<Show Plan>
7	WHILE (1=1) BEGIN SELECT COUNT(*) FRO...	472	5403	0	2015-03-23 20:54:30.073	<Show Plan>

Figure 4-15. Expensive CPU queries

As you can see, the query we used with the OStress tool is the query that got executed the most and took the highest total CPU time. This query could be a good starting point for an investigation. Maybe the query can be optimized or rewritten so it doesn't consume as much CPU time.

SOS_SCHEDULER_YIELD Summary

The SOS_SCHEDULER_YIELD Wait Type will always occur on every SQL Server instance since it is directly related to the scheduling model that SQL Server uses to grant worker threads access to the processor. It can indicate a problem if the total Wait Time or total amount of waiting tasks suddenly increases compared to your baseline measurements. Most of the time a large increase in SOS_SCHEDULER_YIELD waits also means an increase in the CPU load. This increase can either be caused by the SQL Server process itself or by another process outside of SQL Server that requires a large amount of processor time, limiting the time SQL Server can access the processor. If the SQL Server process is responsible for the increase in CPU load you should try to correlate the increase in SOS_SCHEDULER_YIELD waits with an increase in user activity. Another option is to query the sys.dm_exec_query_stats DMV, as shown in Listing 4-1, to find the queries that require the most processor time and focus on optimizing those queries.

THREADPOOL

One of the most notorious Wait Types is the THREADPOOL Wait Type. Unlike the CXPACKET and SOS_SCHEDULER_YIELD Wait Types that occur even if your SQL Server instance isn't experiencing any issues, high THREADPOOL Wait Times do frequently indicate a performance problem. Just as with the other two CPU-related Wait Types we discussed in this book, the THREADPOOL Wait Type is very closely related to the way SQL Server scheduling works.

What Is the THREADPOOL Wait Type?

If you ever see THREADPOOL waits occur on your system with far longer Wait Times than normal, and your SQL Server is (almost) unresponsive, chances are that you are running into an issue called thread pool starvation. Thread pool starvation occurs when there are no more free worker threads available to process requests. When this situation occurs, tasks that are currently waiting to be assigned to a worker thread will log the THREADPOOL Wait Type.

SQL Server provides a number of worker threads to the schedulers with which to process requests. The number of worker threads that are available for your system depends on the number of processors and the processor architecture. The table that follows shows the maximum number of worker threads available for systems with up to 64 logical CPUs.

Table 4-1. Maximum Number of Worker Threads

CPU Number	32-bit Architecture	64-bit Architecture
<= 4	256	512
8	288	576
16	352	704
32	480	960
64	736	1472

You can also calculate the maximum number of worker threads available by using these formulas:

- 32-bit systems with less than, or equal to, 4 logical processors:
256 worker threads
- 32-bit system with more than 4 logical processors:
 $256 + ((\text{number of logical processors} - 4) * 8)$
- 64-bit system with less than, or equal to, 4 logical processors:
512 worker threads
- 64-bit system with more than 4 logical processors:
 $512 + ((\text{number of logical processors} - 4) * 16)$

Even though SQL Server calculates the maximum amount of available worker threads automatically, you can choose to overwrite the default by changing the **Maximum Worker Threads** option inside the Processors properties of your SQL Server instance, as shown in Figure 4-16. By default, the value of the Maximum Worker Threads option will be 0, which means SQL Server will calculate and assign the maximum amount of worker threads available using the formulas above.

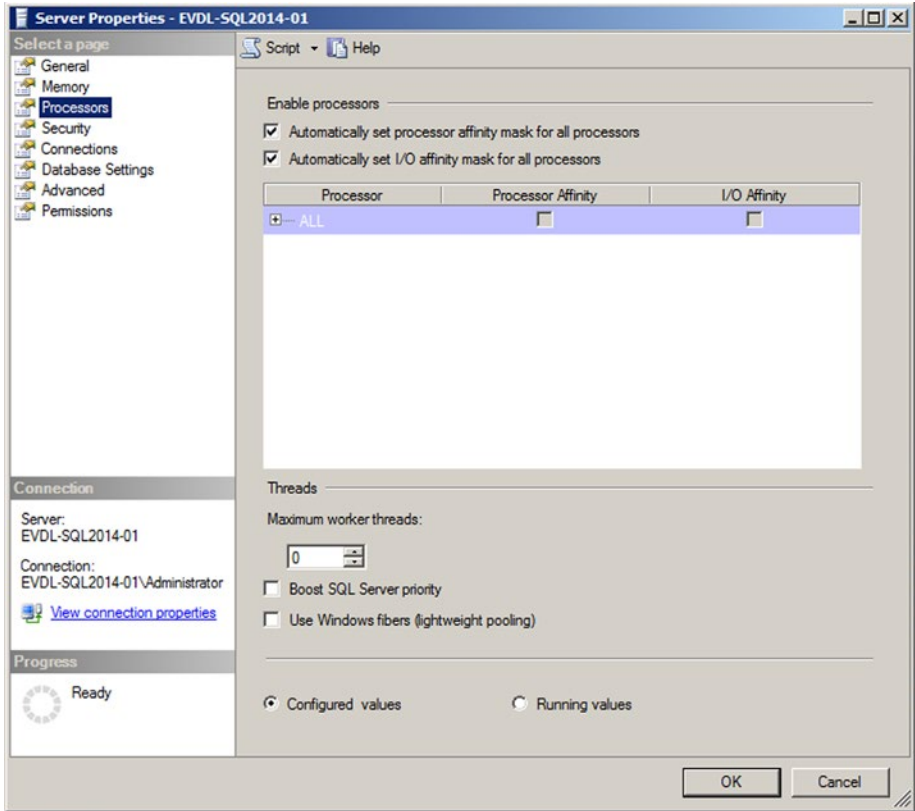


Figure 4-16. Processors configuration of a SQL Server instance

You can also query the number of worker threads assigned to your SQL Server instance by running the following query:

```
SELECT
    max_workers_count
FROM sys.dm_os_sys_info;
```

For my 64-bit test SQL Server that has four logical processors, I have 512 worker threads available, as you can see in Figure 4-17.

	max_workers_count
1	512

Figure 4-17. Amount of worker threads on my test machine

One piece of advice I frequently read on the Internet related to THREADPOOL waits is to change the Maximum Worker Threads option to a value higher than the one your SQL Server instance has by default. I strongly advise against changing this option from its default value. Changing the setting to a higher value than the amount of worker threads you would receive by default can actually degrade the performance of your SQL Server because context-switching occurs far more often. Another reason not to change the setting is that every worker thread requires a bit of memory to operate; for 32-bit systems this is 512 KB per worker thread, and for 64-bit systems it's 2048 KB.

THREADPOOL Example

Let's start with an example of THREADPOOL waits occurring on my test SQL Server instance. Even though I have warned you multiple times already about making sure to not run any of the demo scripts in this book on a production environment, this one deserves a special reminder. Running the demo scripts in this section can cause your SQL Server to become completely unresponsive, not accepting any new connections, and can eventually require a restart of the SQL Server service! Do not run this on a SQL Server that isn't allowed to become unresponsive!

For this example we are going to use the Ostress utility again to simulate concurrency and load against the test SQL Server instance. First, we create another .sql file (`select_rnd.sql`) that holds the following query that we will execute using Ostress:

```
SELECT TOP 1 *
FROM Sales.SalesOrderDetail
ORDER BY NEWID() OPTION (MAXDOP 1)
```

This query will select one random row from the `Sales.SalesOrderDetail` table in the AdventureWorks database. There is a reason I included the query option to serially run this query, and I will explain it later on.

Now, before we launch Ostress to execute the query above, we are purposely going to lower the maximum amount of worker threads available on the test SQL Server. To do this we execute this query:

```
EXEC sp_configure 'show advanced options', 1;
GO
RECONFIGURE
GO
EXEC sp_configure 'max worker threads', 128;
GO
RECONFIGURE
GO
```

This will set the maximum number of worker threads available to 128, the minimum value for a 64-bit SQL Server instance.

Let's fire up Ostress and execute the .sql script we created earlier:

```
"C:\Program Files\Microsoft Corporation\RMLUtils\ostress.exe" -E
-dAdventureWorks -i"C:\select_rnd.sql" -n150 -r10 -q
```

In this case we will start 150 different threads that will execute the query in the `select_rnd.sql` file 10 times. The reason for spawning 150 threads is that this value is higher than the maximum amount of worker threads available on the test SQL Server instance, but not so high that we cannot execute queries anymore.

While the script is running, let's take a look at the number of worker threads running and waiting using the `sys.dm_os_schedulers` DMV:

```
SELECT
    scheduler_id,
    current_tasks_count,
    runnable_tasks_count,
    current_workers_count,
    active_workers_count,
    work_queue_count
FROM sys.dm_os_schedulers
WHERE status = 'VISIBLE ONLINE';
```

The results of this query are shown in Figure 4-18.

	scheduler_id	current_tasks_count	runnable_tasks_count	current_workers_count	active_workers_count	work_queue_count
1	0	63	49	56	56	7
2	1	35	21	33	31	2
3	2	34	20	28	28	6
4	3	53	30	39	38	14

Figure 4-18. Tasks and worker threads per scheduler

The most important columns here are the `current_workers_count`, `active_workers_count`, and `work_queue_count` columns. The `current_workers_count` column shows the number of worker threads associated with this scheduler; this number also includes worker threads that are not yet assigned to a task. The `active_workers_count` column returns the number of worker threads that are in the “RUNNING,” “RUNNABLE,” or “SUSPENDED” states. The big difference between the `current_workers_count` and the `active_workers_count` columns is that the `active_workers_count` is the number of worker threads that have been assigned to a task, while the `current_workers_count` returns all the worker threads. The `work_queue_count` column shows us the number of tasks that are currently waiting to get a worker thread assigned to them. If you see values higher than 0 in this column for a longer period of time and for all schedulers, you are experiencing thread pool starvation.

Let's check the `sys.dm_os_waiting_tasks` DMV for waiting tasks that originate from a user session. Notice that we filter out all the sessions that have a session ID lower than 50, even though I told you to not do this in Chapter 2:

```
SELECT *
FROM sys.dm_os_waiting_tasks
WHERE session_id > 50;
```

If we check the results on the test SQL Server instance, we could conclude that nothing is waiting, as you can see in Figure 4-19. The test SQL Server is responding incredibly slowly though, and querying anything requires multiple seconds.

waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address	blocking_task_address

Figure 4-19. No tasks are waiting

Let's check the `sys.dm_os_waiting_tasks` DMV without filtering out session IDs:

```
SELECT *
FROM sys.dm_os_waiting_tasks;
```

As you can see in Figure 4-20, `THREADPOOL` waits are not logged as user sessions, but actually have an empty session ID. This is the reason I always recommend to not filter the `sys.dm_os_waiting_tasks` DMV on session ID numbers.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type
14	0x000000016DE84558	NULL	NULL	34969	THREADPOOL
15	0x000000016EF8A558	NULL	NULL	34969	THREADPOOL
16	0x000000016EF8B868	NULL	NULL	34969	THREADPOOL
17	0x000000016EF8ACF8	NULL	NULL	29851	THREADPOOL
18	0x00000001720D0928	NULL	NULL	28612	THREADPOOL
19	0x00000001720D1498	NULL	NULL	27724	THREADPOOL
20	0x000000016DE85C38	NULL	NULL	27289	THREADPOOL

Figure 4-20. `THREADPOOL` waits

There are quite a lot of `THREADPOOL` waits, and they all have very high Wait Times of between 20 and 40 seconds. Things can get even worse than this though. Figure 4-21 shows an error I encountered when I tried to connect to my test SQL Server instance while running the `Ostrestool`.

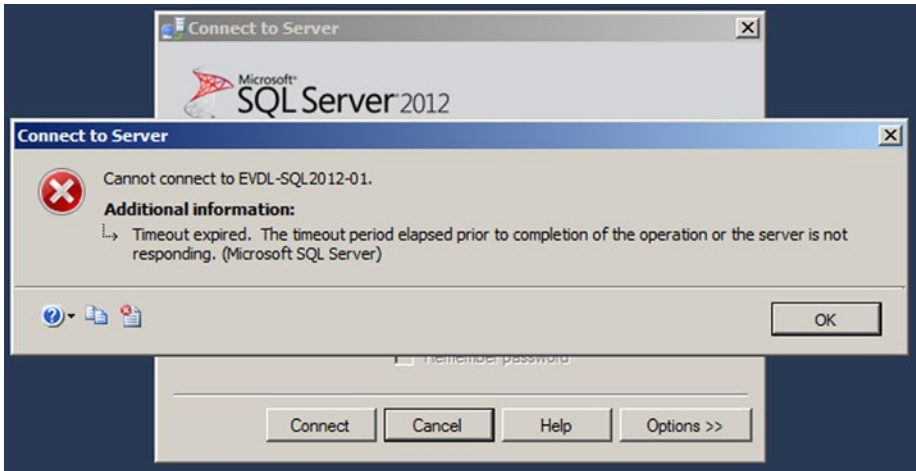


Figure 4-21. Timeouts are occurring, SQL Server is unresponsive

Now that we have seen the kind of problems thread pool starvation can create, let's take a look at how we can lower, or even resolve, THREADPOOL waits.

Gaining Access to Our SQL Server during THREADPOOL Waits

THREADPOOL waits can be very difficult to troubleshoot, mostly because there are many possible reasons why your SQL Server doesn't have any free worker threads available. Also, THREADPOOL waits can completely lock down your SQL Server instance, making connections to it (and troubleshooting it) almost impossible, as you have seen in the earlier example.

The first step you should take to make sure you do not get into a situation where you cannot connect to your SQL Server instance for troubleshooting is to enable the Dedicated Administrator Connection (or DAC). If you remember the section about schedulers in Chapter 1, you might recall a special type of scheduler reserved for the DAC. This dedicated scheduler, shown in Figure 4-22, is strictly reserved for the DAC and has access to its own worker threads.

	scheduler_address	parent_node_id	scheduler_id	cpu_id	status	is_online	is_idle
1	0x000000017F140040	0	0	0	VISIBLE ONLINE	1	1
2	0x000000017F160040	0	1	1	VISIBLE ONLINE	1	1
3	0x000000017F280040	0	2	2	VISIBLE ONLINE	1	0
4	0x000000017F2A0040	0	3	3	VISIBLE ONLINE	1	1
5	0x000000017F2C0040	0	1048578	0	HIDDEN ONLINE	1	0
6	0x000000017F8C0040	64	1048576	0	VISIBLE ONLINE (DAC)	1	1
7	0x000000017D8E0040	0	1048579	1	HIDDEN ONLINE	1	1

Figure 4-22. Dedicated Administrator Connection scheduler

If you connect through the DAC to your SQL Server instance, your session will be mapped to the DAC scheduler. This makes it possible to connect and execute queries even if all the other schedulers have massive task queues.

You can enable the DAC by executing the following query:

```
sp_configure 'remote admin connections', 1
GO
RECONFIGURE
GO
```

If you want to connect to your SQL Server instance using the DAC you need to add the ADMIN: prefix to the server name you are connecting to, as shown in Figure 4-23. You can only connect using the DAC when you execute a new query from inside SQL Server Management Studio without being connected to the server.



Figure 4-23. Connect using the Dedicated Administrator Connection

Now that you are able to connect to your SQL Server instances using the DAC, you always have a way in, even when the SQL Server instance won't accept any new connections.

With the DAC enabled, let's discuss some common causes for THREADPOOL waits.

Lowering THREADPOOL Waits Caused by Parallelism

One of the most common causes for THREADPOOL waits I encounter is related to the extensive use of parallelism during query execution. During the execution of a parallel query, multiple worker threads are used to perform the work needed. If you left the configuration options

related to parallelism—Max Degree of Parallelism and Cost Threshold of Parallelism—at the default values, it might cause more queries to run in parallel than was intended. Depending on how many processors your SQL Server has access to, and the number of worker threads used during a parallel query, one single parallel query can require many worker threads.

If you run into this specific case of high and frequent THREADPOOL waits you will usually see many CXPACKET waits as well (sometimes with high Wait Times). To show this behavior I have modified the query we used to generate THREADPOOL waits so that it will execute using parallelism. In this case I commented out the MAXDOP query option:

```
SELECT TOP 1 *
FROM Sales.SalesOrderDetail
ORDER BY NEWID()
-- OPTION (MAXDOP 1)
```

For this example I also configured the Max Degree of Parallelism to its default value of 0, and set the Cost Threshold for Parallelism option to 1. This way I am 100% sure the query will be run using parallelism. I left the Max Worker Threads option on a value of 128 as we configured earlier.

If we now repeat the same Ostress test we performed earlier in this chapter by executing the command below, we should see THREADPOOL waits occur again in the `sys.dm_os_waiting_tasks` DMV, as shown in Figure 4-24.

```
"C:\Program Files\Microsoft Corporation\RMLUtils\ostress.exe" -E
-dAdventureWorks -i"C:\select_rnd.sql" -n150 -r10 -q
```

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address
1	0x0000000174062928	NULL	NULL	4966	THREADPOOL	NULL
2	0x0000000172427498	NULL	NULL	6382	THREADPOOL	NULL
3	0x0000000172427C38	NULL	NULL	6382	THREADPOOL	NULL
4	0x0000000172426558	NULL	NULL	6090	THREADPOOL	NULL
5	0x0000000172C93C38	NULL	NULL	4616	THREADPOOL	NULL
6	0x0000000165A78558	NULL	NULL	3959	THREADPOOL	NULL
7	0x00000001657F4558	NULL	NULL	1568	THREADPOOL	NULL
8	0x0000000172C8C558	NULL	NULL	1520	THREADPOOL	NULL

Figure 4-24. THREADPOOL waits

But this time, because our test query is executed in parallel, we will also find many CXPACKET waits returned by the `sys.dm_os_waiting_tasks` DMV, as shown in Figure 4-25.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address
8	0x0000000172C8C188	111	0	2135	CXPACKET	0x0000000173D9F950
9	0x0000000173728928	114	0	2002	CXPACKET	0x000000016F0998E0
10	0x0000000172C8D868	116	0	1921	CXPACKET	0x0000000161FEB950
11	0x00000001692C70C8	115	0	1361	CXPACKET	0x0000000165861900
12	0x00000001741CB0C8	117	0	739	CXPACKET	0x000000016B9A78E0
13	0x00000001692C6188	124	0	468	CXPACKET	0x000000016043D950
14	0x00000001692C7C38	NULL	NULL	1498	THREADPOOL	NULL
15	0x0000000173728188	NULL	NULL	1474	THREADPOOL	NULL
16	0x00000001692C6CF8	NULL	NULL	1370	THREADPOOL	NULL
17	0x00000001692C6928	NULL	NULL	1366	THREADPOOL	NULL

Figure 4-25. CXPACKET and THREADPOOL waits

If you see this behavior occurring on your SQL Server instance it might be worth the effort to check your parallelism configuration. The first section of this chapter discussed CXPACKET waits and how you can lower them. Another hint that might steer you in this direction is that the CPU load during this particular case is usually higher than normal. In the case of my test SQL Server instance, all my CPUs went to 100%.

Lowering THREADPOOL Waits Caused by User Connections

Another common cause of THREADPOOL waits is a sudden increase in the number of users connecting and executing queries against your SQL Server instance. This problem can occur if, for instance, the application that is connecting to your SQL Server instance uses multiple connections. The main problem here is that those connections stay active and keep acquiring worker threads.

To give you an example of this problem we will again use Ostress to connect and execute queries against my test SQL Server instance. In this case we will use a different .sql file, saved as wait.sql, as input for Ostress, with the following query inside it:

```
WAITFOR DELAY '00:05:00'
```

The only thing this query will do is wait for five minutes. After those five minutes the query will end and the connection will disconnect.

Let's run Ostress using the wait.sql file:

```
"C:\Program Files\Microsoft Corporation\RMLUtils\ostress.exe" -E
-dAdventureWorks -i"C:\wait.sql" -n120 -r1 -q
```

We change the number of threads generated by Ostress to 120 and again leave the Max Worker Threads option set to 128 worker threads.

When we query the `sys.dm_exec_sessions` DMV using the query below, we see that many new user sessions, generated by the `Ostress` utility, are active, as shown in Figure 4-26.

```
SELECT *
FROM sys.dm_exec_sessions
WHERE is_user_process = 1;
```

	session_id	login_time	host_name	program_name
7	57	2015-03-26 11:25:15.530	EVDL-SQL2012-01	OSTRESS
8	58	2015-03-26 11:25:15.533	EVDL-SQL2012-01	OSTRESS
9	59	2015-03-26 11:25:15.540	EVDL-SQL2012-01	OSTRESS
10	60	2015-03-26 11:25:15.547	EVDL-SQL2012-01	OSTRESS
11	61	2015-03-26 11:25:15.547	EVDL-SQL2012-01	OSTRESS
12	62	2015-03-26 11:25:15.547	EVDL-SQL2012-01	OSTRESS
13	63	2015-03-26 11:25:15.547	EVDL-SQL2012-01	OSTRESS

Figure 4-26. *Ostress user sessions*

If we query the `sys.dm_os_waiting_tasks` DMV we see that `THREADPOOL` waits are occurring, as shown in Figure 4-27.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address
1	0x000000017F01D868	NULL	NULL	11264	THREADPOOL	NULL
2	0x0000000167610CF8	NULL	NULL	11315	THREADPOOL	NULL
3	0x0000000167611868	NULL	NULL	11315	THREADPOOL	NULL
4	0x00000001676110C8	NULL	NULL	11313	THREADPOOL	NULL
5	0x0000000167610558	NULL	NULL	11310	THREADPOOL	NULL
6	0x00000001741B3C38	NULL	NULL	11306	THREADPOOL	NULL
7	0x00000001741B2558	NULL	NULL	11304	THREADPOOL	NULL
8	0x00000001741B3868	NULL	NULL	11296	THREADPOOL	NULL

Figure 4-27. *THREADPOOL waits inside the sys.dm_os_waiting_tasks DMV*

The big difference between `THREADPOOL` waits caused by excessive parallelism and an increase in user connections is that the CPU of my test SQL Server instance remains low in the latter case, as shown in Figure 4-28.




Figure 4-28. *CPU usage*

The small spike in the CPU Usage History graph is caused by starting up the Ostress utility. After that, the CPUs remain at a constant low usage percentage.

Resolving THREADPOOL waits caused by an increase in user connections should start at the source. Where do the user connections come from? What are those connections performing? I have seen cases where an application suddenly used hundreds of active user connections after an update, and as the SQL Server instance was not designed to handle that amount of concurrent, active, connections, THREADPOOL waits appeared.

Keep in mind that the user connections should only cause THREADPOOL waits when they are actually running queries. User connections that are connected to the SQL Server instance but are not executing anything should not be a reason for THREADPOOL waits.

Also, having many different user connections be active against a database can create many locks on rows or tables. If you notice high lock-related Wait Times together with THREADPOOL waits, the problem could be the high amount of locking and blocking occurring. In this case you should try to find the queries that are causing the lock waits and see if you can optimize them. We will discuss lock-related Wait Types, and what you can do about them, in Chapter 7, “Lock-Related Wait Types.”

THREADPOOL Summary

THREADPOOL waits are one of the most alarming Wait Types to see on your SQL Server instance. They occur because there are not enough free worker threads available to process requests, so tasks that request a worker thread will have to wait until a new worker thread becomes available. Thankfully, THREADPOOL waits are not very common, as they have the potential to completely lock you out of your SQL Server instance. The only way to connect in those cases is by using the Dedicated Administrator Connection (or DAC), which I urge you to enable on all your SQL Server instances.

Excessive use of parallelism and a large increase in active user connections are two of the most common causes for THREADPOOL waits. The former has a direct relation to the CXPACKET Wait Type we discussed earlier, so methods to resolve the CXPACKET Wait Type can also help to resolve THREADPOOL waits. The latter requires a deeper investigation into why the number of active user connections suddenly increased. Maybe they are the result of a bug in the application connection to the SQL Server instance. We also briefly touched on locking and blocking behavior as a possible cause for THREADPOOL waits. We will take a deeper look at how we can resolve lock-related waits in Chapter 7, “Lock-Related Wait Types.”

CHAPTER 5



IO-Related Wait Types

In this chapter we will take a look at IO-related Wait Types in the broadest sense of the term. I selected Wait Types that are related to the storage, memory, or network components of your system. One could argue that the majority of Wait Types will fit this category, and that is probably right, but to prevent this chapter from covering 90% of all the Wait Types in this book, I had to choose carefully. I consider the Wait Types in this chapter to have a direct relation to storage, memory, or network but to not relate directly to a functionality or concept in SQL Server. For instance, the PAGEIOLATCH_XX Wait Types are frequently related to storage, but they are not included in this chapter. The reason for this is because they are also a latch Wait Type, and I believe latch Wait Types deserve a separate chapter because of their function in SQL Server.

The performance of the IO-related components are incredibly important for SQL Server. Practically every part of SQL Server interacts with these components in one way or another, whether it is a data page that needs to be read from disk into memory or the results from a query that need to be transported across the network to your end users. If one of these components can't handle the workload you are generating on your SQL Server instance, or isn't configured properly, your performance will decline.

The Wait Types in this chapter can help you track down which of your IO components is slowing you down so you can take appropriate action to prevent or resolve performance-related incidents.

ASYNC_IO_COMPLETION

The ASYNC_IO_COMPLETION Wait Type is a pretty common Wait Type that occurs every time SQL Server performs a file-related action on the storage subsystem and has to wait for it to complete. You will frequently see this Wait Type when you are performing actions that interact with the storage subsystem, like a backup. Just like with most Wait Types, if you are seeing this Wait Type occur it doesn't necessarily mean there is a problem with your storage subsystem. It will only become a problem if the Wait Time is longer than you expect it to be compared to your baseline values, which we discussed in Chapter 3, "Building a Solid Baseline."

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-1-4842-1139-7_5](https://doi.org/10.1007/978-1-4842-1139-7_5)) contains supplementary material, which is available to authorized users.

What Is the ASYNC_IO_COMPLETION Wait Type?

If we look up the `ASYNC_IO_COMPLETION` Wait Type in Books Online (BOL) we will see the following definition: “Occurs when a task is waiting for I/Os to finish.” This is a rather short and vague definition. Let’s add a little more detail. `ASYNC_IO_COMPLETION` waits occur when a task is waiting for a storage-related action to finish. The task is initiated and monitored by SQL Server. Figure 5-1 shows this as a visual representation of the Wait Type.




Figure 5-1. `ASYNC_IO_COMPLETION` wait occurring

As long as the storage-related action is running, the `ASYNC_IO_COMPLETION` Wait Time is being logged. As you can imagine, the faster your storage subsystem, the lower your `ASYNC_IO_COMPLETION` Wait Times will be.

As I said earlier, usually `ASYNC_IO_COMPLETION` waits are no cause for concern. They will happen normally during many SQL Server operations that need to access the storage subsystem, like backups or the creation of a new database. It can become a cause for concern if the Wait Times are higher than you expected when compared to your baseline measurements.

ASYNC_IO_COMPLETION Example

Let’s go through an example that generates `ASYNC_IO_COMPLETION` waits. We won’t need any extra utilities for this; just running a database backup will trigger `ASYNC_IO_COMPLETION` waits.

In this case I will perform a backup of the AdventureWorks database on my test server.

To perform this action I will use the query in Listing 5-1. This query will reset the `sys.dm_os_wait_stats` DMV, perform the database backup, and then query the `sys.dm_os_wait_stats` DMV for the `ASYNC_IO_COMPLETION` waits.

Listing 5-1. Generate `ASYNC_IO_COMPLETION` waits

```
USE [master]
GO

DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR);

BACKUP DATABASE [AdventureWorks]
  TO DISK = N'F:\Backup\aw_backup.bak'
```

```

WITH
    NAME = N'AdventureWorks-Full Database Backup',
    STATS = 2;
GO

SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'ASYNC_IO_COMPLETION';

```

The backup operation took 11 seconds on my test SQL Server instance. Figure 5-2 shows the results of the query in Listing 5-1.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	ASYNC_IO_COMPLETION	2	10592	10589	0

Figure 5-2. ASYNC_IO_COMPLETION Wait Time

As you can see, for almost the entire duration of the database backup ASYNC_IO_COMPLETION waits were logged.

Lowering ASYNC_IO_COMPLETION Waits

One common cause of high ASYNC_IO_COMPLETION Wait Times is a database backup, as you just saw in the example. If you want to find out if your ASYNC_IO_COMPLETION waits are occurring because a backup is being performed, try to look for backup-related waits occurring at the same time.

If we were to slightly modify the last `sys.dm_os_waiting_tasks` query in Listing 5-1, we would see backup-related Wait Types being returned by the DMV:

```

SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type IN
    (
        'ASYNC_IO_COMPLETION',
        'BACKUIO',
        'BACKUPBUFFER'
    );

```

Figure 5-3 shows the result of the query in Listing 5-1, but with this modification.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	ASYNC_IO_COMPLETION	2	11867	11861	0
2	BACKUPBUFFER	93	3924	256	244
3	BACKUIO	225	9050	552	25

Figure 5-3. ASYNC_IO_COMPLETION waits together with backup-related waits

If you see both occurring at the same time, chances are that a database backup is causing your `ASYNC_IO_COMPLETION` waits.

Another possible method of lowering `ASYNC_IO_COMPLETION` waits is by configuring instant file initialization. Instant file initialization was introduced in Windows 2003 and speeds up the process of allocating space on a disk tremendously by removing the need to zero-out files (writing zeros inside files before they can get used). This does not affect the speed of your backup, but will give increased performance when creating a database, adding files to a database, or restoring a database. Instant file initialization is not enabled by default, unless you are running your SQL Server service under an account that has local administrator privileges.

To make use of instant file initialization we need to change a local policy on the machine SQL Server is running on and add the account your SQL Server service is running under.

You can find this policy by opening the Local Security Policy MMC under Administrative Tools in the Configuration Panel. Open up the Local Policies - User Rights Assignment folder and scroll down to the “Perform volume maintenance tasks” policy as shown in Figure 5-4.

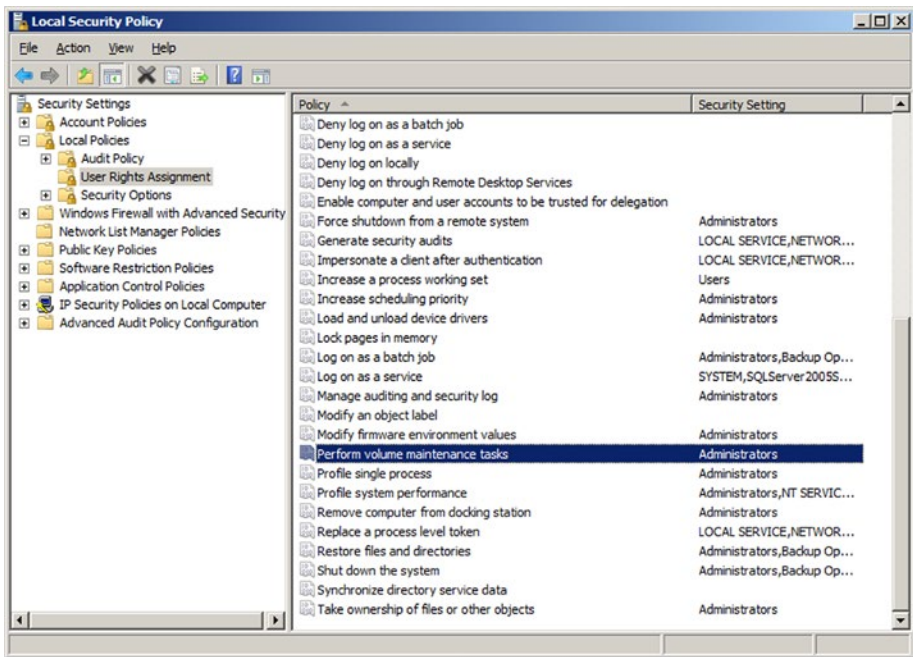


Figure 5-4. Perform volume maintenance tasks local policy

Double-click on the policy to open it and add the account your SQL Server service is running under. The last step is restarting your SQL Server service. After the restart, SQL Server can make use of instant file initialization.

To show you the impact of instant file initialization, I used the query in Listing 5-2. This query clears the `sys.dm_os_wait_stats` DMV, then creates a new database with a 500 MB data file and a 100 MB log file. It then queries the `sys.dm_os_wait_stats` DMV for the `ASYNC_IO_COMPLETION` Wait Type.

Listing 5-2. Measure the impact of instant file initialization on `ASYNC_IO_COMPLETION` waits

```
DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR);

CREATE DATABASE [IO_test]
ON PRIMARY
(
    NAME = N'IO_test', FILENAME = N'E:\Data\IO_test.mdf' , SIZE = 512000KB ,
    FILEGROWTH = 10%
)
LOG ON
(
    NAME = N'IO_test_log', FILENAME = N'E:\Log\IO_test_log.ldf' ,
    SIZE = 102400KB , FILEGROWTH = 10%
);
GO

SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'ASYNC_IO_COMPLETION';
```

Figure 5-5 shows the Wait Statistics information both before and after configuring instant file initialization.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	ASYNC_IO_COMPLETION	1	10591	10591	0

before

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	ASYNC_IO_COMPLETION	1	1878	1878	0

after

Figure 5-5. Impact of instant file initialization on `ASYNC_IO_COMPLETION` waits

Even for this relatively small database the gain of using instant file initialization is pretty big, as you can see in the difference in Wait Times. Before enabling instant file initialization, the query in Listing 5-2 took 11 seconds to complete; after the change it went down to 2 seconds.

If you configured instant file initialization and checked that no backups are being performed at the same time that you are seeing high `ASYNC_IO_COMPLETION` waits, the problem might be your storage subsystem. A good method of analyzing potential storage problems is by using Perfmon to monitor the `Avg. Disk/sec Read` and `Avg. Disk/sec Write` counters on the disks on which your database resides, as shown in Figure 5-6. These counters show you the read and write latency to your disks in seconds (this means a value of 0.005 means 5 milliseconds). SQL Server performs optimally with a maximum latency of 5 milliseconds. Above 20 milliseconds, latency will cause noticeable performance degradation. The higher the latency value, the higher the Wait Time of storage-related Wait Types will be.

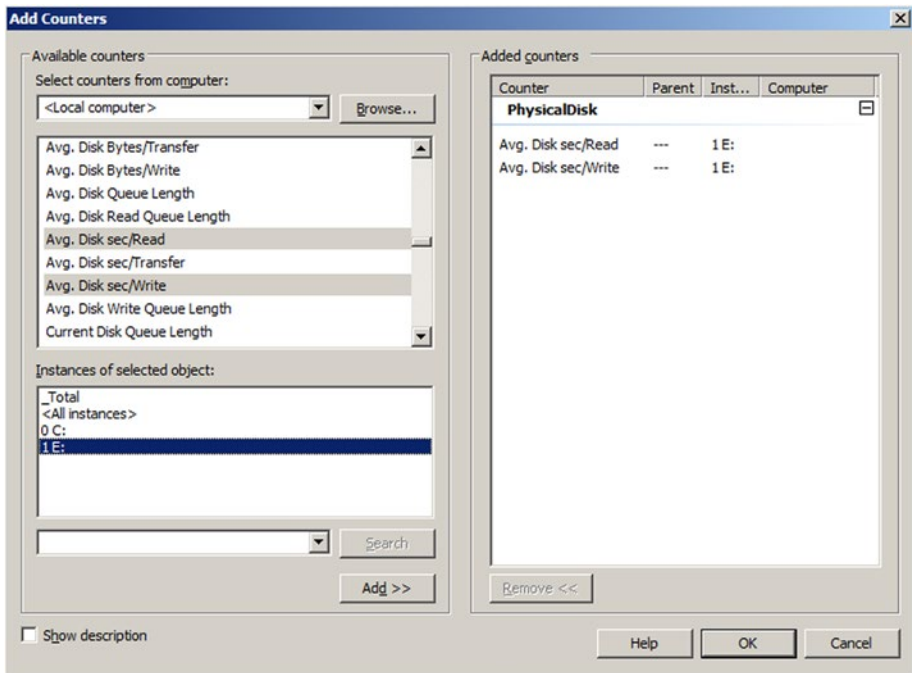


Figure 5-6. *Avg. Disk sec/Read and Avg. Disk sec/Write Perfmon counters*

Be careful about jumping to conclusions regarding your storage performance. Always talk to your storage administrator (if you have one) and show him/her your measurements before you decide the storage subsystem is the bottleneck. Storage is the domain of the storage administrator, and he/she can help you analyze and resolve performance problems.

ASYNC_IO_COMPLETION Summary

The `ASYNC_IO_COMPLETION` Wait Type occurs when you perform actions related to the storage subsystem from inside your SQL Server instance, most notably database backups and the creation of new databases. While `ASYNC_IO_COMPLETION` waits are completely normal, they can indicate storage-related problems if Wait Times are higher than normal. Before you run to your storage administrator, make sure there is actually a performance problem. One possible way to do this is by checking your storage latency, as high latency values will impact `ASYNC_IO_COMPLETION` Wait Times as well. Also check whether the higher `ASYNC_IO_COMPLETION` Wait Times are directly related to database backups being performed. One great method to lower `ASYNC_IO_COMPLETION` Wait Times is by enabling instant file initialization by adding your SQL Server service account to the `Perfmom` volume maintenance tasks local policy.

ASYNC_NETWORK_IO

Just like the `ASYNC_IO_COMPLETION` Wait Type, the `ASYNC_NETWORK_IO` Wait Type is related to throughput. But instead of storage-subsystem throughput, the `ASYNC_NETWORK_IO` Wait Type is related to the throughput of your network connection between your SQL Server instance and your clients. Again, seeing Wait Times for this specific Wait Type does not necessarily mean there is a network-related issue, since `ASYNC_NETWORK_IO` waits always occur, even if you query your SQL Server instance on the SQL Server itself.

What Is the `ASYNC_NETWORK_IO` Wait Type?

`ASYNC_NETWORK_IO` waits usually occur when client applications cannot process the query results fast enough, or when you have a network-related performance problem. The former will in most cases be the most likely, since many applications process SQL Server results on a row-by-row basis, or simply cannot handle the amount of data. This forces the SQL Server to wait on sending query results across the network. While SQL Server is waiting to send the requested data, the `ASYNC_NETWORK_IO` Wait Type is logged. Another situation in which `ASYNC_NETWORK_IO` waits can occur is when you are using a linked server to query remote databases. Figure 5-7 shows a graphical representation of this.




Figure 5-7. `ASYNC_NETWORK_IO`

ASYNC_NETWORK_IO Example

Showing an example of the ASYNC_NETWORK_IO Wait Type doesn't require a complicated test environment. Listing 5-3 shows a query that will generate ASYNC_NETWORK_IO waits when run against my test SQL Server instance from another computer using SQL Server Management Studio, which should be enough. The query is going to clear the `sys.dm_os_wait_stats` DMV then perform the actual query against the AdventureWorks database. The last statement will show us the Wait Times of the ASYNC_NETWORK_IO Wait Type.

Listing 5-3. Generate ASYNC_NETWORK_IO waits

```
DBC SQLPERF ('sys.dm_os_wait_stats', CLEAR);

SELECT *
FROM Person.Person;

SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'ASYNC_NETWORK_IO';
```

Figure 5-8 shows the Wait Times for the ASYNC_NETWORK_IO Wait Type.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	ASYNC_NETWORK_IO	165	516	58	50

Figure 5-8. ASYNC_NETWORK_IO Wait Times

In this example, the results of the query against the AdventureWorks database couldn't be processed by the SQL Server Management Studio application as fast as the SQL Server instance supplied the results, and ASYNC_NETWORK_IO waits occurred.

Lowering ASYNC_NETWORK_IO Waits

One of the "easiest" ways to lower ASYNC_NETWORK_IO waits is to identify queries that will return a large result set back to the application. For instance, if we modify the query to return only the first 100 rows, the SQL Server Management Studio might be able to keep up with the information returned to it:

```
DBC SQLPERF ('sys.dm_os_wait_stats', CLEAR);

SELECT TOP 100 *
FROM Person.Person;

SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'ASYNC_NETWORK_IO';
```

The resulting Wait Times after this modification can be seen in Figure 5-9.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	ASYNC_NETWORK_IO	0	0	0	0

Figure 5-9. ASYNC_NETWORK_IO Wait Times after modifying the query

As you can see, we didn't run into any ASYNC_NETWORK_IO waits this time. The SQL Server Management Studio was able to keep up with the results returned, so the SQL Server instance we queried didn't have to delay sending the results back to the client.

Another way to limit results returned could be by filtering out information using WHERE clauses that isn't used by the application in the first place. Smaller results will result in lower ASYNC_NETWORK_IO Wait Times.

If you believe that ASYNC_NETWORK_IO waits are not caused by large results being returned to an application, or by the speed at which an application can process the results, there is also a possibility that your network configuration is slowing you down. In this case you should first check your network utilization. Sadly there isn't a counter in Perfmon that directly shows the network utilization without having you perform some math to calculate it. Instead, you can use the Networking tab of the Task Manager to view your network-card utilization, as shown in Figure 5-10.




Figure 5-10. Task Manager Network utilization

If you notice that the network utilization is high while you are experiencing higher than normal `ASYNC_NETWORK_IO` Wait Times, it could be possible that the network is slowing you down. In that case it might be a good idea to talk to your network administrator. Usually a network configuration consists of many parts, like switches, routers, firewalls, network cables, drivers, firmware, potential virtualization of the operating system, and so on. All of these parts can slow down your network throughput and can be a potential cause of `ASYNC_NETWORK_IO` waits.

ASYNC_NETWORK_IO Summary

The `ASYNC_NETWORK_IO` Wait Type occurs whenever an application requests query results from a SQL Server instance over the network and cannot process the returned results fast enough. Seeing `ASYNC_NETWORK_IO` waits occur is completely normal, but higher than normal Wait Times can be caused by changes in the returned query results or network-related problems. Lowering `ASYNC_NETWORK_IO` Wait Times that are application related can be achieved by decreasing the number of rows and/or columns returned to the application.

CMEMTHREAD

Waits of the `CMEMTHREAD` Wait Type are memory related and indicate a pressure on certain SQL Server–related memory objects. These memory objects allocate memory for the various parts of SQL Server like the buffer cache and the procedure cache. Whenever `CMEMTHREAD` waits occur it means that multiple threads are trying to access the same memory object at the same time.

What Is the CMEMTHREAD Wait Type?

To explain how `CMEMTHREAD` Wait Type generation works, we have to dig a little deeper inside some programming terminology, specifically the terms *mutual exclusions*, *critical sections*, and *thread safety*. These three concepts play a direct role in the `CMEMTHREAD` Wait Type generation.

A critical section consists of a piece of code that accesses a shared resource that can only be accessed by one thread at a time. In our case the shared resource would be a SQL Server memory object. The SQL Server memory objects can only be accessed one thread at a time so as to ensure no corruption to the memory object can occur. Because there are many threads that want access to memory objects, we have to use a method to ensure only one thread gets access at a time. This method is called mutual exclusion. SQL Server uses a Mutex object to make sure concurrent threads are not in their critical sections at the same time when accessing the memory object. A Mutex does this by serializing the thread access to the memory object. Only a single thread can be the owner of a Mutex object, and while a thread has ownership it can access the shared resource. When the thread is done, the Mutex object will move to the next thread in line. By using these objects we have created thread-safe code, where multiple threads do not have concurrent access to memory objects. Figure 5-11 shows this situation.




Figure 5-11. Thread waiting for a Mutex object to access a shared resource

A simplified example of this behavior is when you and a large group of other people are waiting for a single ticket dispenser to buy a ticket to see your favorite rock band. In this case the ticket dispenser is the shared resource, and only a single person can access the ticket dispenser at a time. When we reach the ticket dispenser we can get a ticket, and the people behind us have to wait till it's their time. After we have bought a ticket the next person in line gets access to the ticket dispenser.

We can also view this behavior in SQL Server, but to do this we have to make use of a debugger (like WinDbg). Figure 5-12 shows how a thread in SQL server waits for a Mutex since access to the memory object was granted to another thread. To capture this image I used an Extended Event session that created a SQL Server mini-dump when a CMEMTHREAD wait occurred. I then used WinDbg to open the mini-dump and returned the stack.

```
0:170> ln 0x000007FEF9F86059:ln 0x000007FEF9F4219C:ln 0x000007FEF9F41F03:ln 0x000007FEF9F437AB:
(000007fe`f9f41400) sqldk!XeSocPkg::wait_info::Publish+0x138 | (000007fe`fa020930) sqldk
(000007fe`f9f41f70) sqldk!SOS_Scheduler::UpdateWaitTimeStats+0x2bc | (000007fe`f9f42250)
(000007fe`f9f41e90) sqldk!SOS_Task::PostWait+0x9e | (000007fe`f9f4db60) sqldk!SOS_Task::
(000007fe`f9f437e0) sqldk!EventInternal<SuspendQueueSLock>::Wait+0x2ca | (000007fe`f9f43a6
(000007fe`f9f6e1a0) sqldk!SOS_UnfairMutexPair::LongWait+0x191 | (000007fe`f9f6e400) sqld
(000007fe`f9f455d0) sqldk!SOS_UnfairMutexPair::AcquirePair+0x46 | (000007fe`f9f45640) sq
(000007fe`f9f45970) sqldk!CMemThread<CMemObj>::Alloc+0xb6 | (000007fe`f9f4ee40) sqldk!CM
```

Figure 5-12. Example of a CMEMTHREAD wait occurring in a mini-dump

The important line here is the `SOS_UnfairMutexPair::LongWait`, which generates the CMEMTHREAD wait because the thread we are monitoring here has to wait for another thread that currently has access to the memory object. The line after that, `SOS_UnfairMutexPair::AcquirePair`, means the thread received the Mutex, followed by access to the memory object represented by `CMemThread<CMemObj>::Alloc`.

Lowering CMEMTHREAD Waits

Since there are many different memory objects present in SQL Server that could potentially generate CMEMTHREAD waits, there are many possible solutions to lowering CMEMTHREAD Wait Times depending on the memory object that is being accessed.

One of the more common situations where CMEMTHREAD waits can occur is when large amounts of short, concurrent, ad-hoc queries are being executed. Every time an ad-hoc query is executed that could not be parameterized, the Query Optimizer will generate a new execution plan for the query. All these new execution plans need to be entered into the procedure cache, and a memory object for allocating cache descriptors is accessed.

Since the memory object is thread-safe, CMEMTHREAD waits can occur if the rate of insertion is high enough. A good place to start looking if you suspect CMEMTHREAD waits are going to occur because of ad-hoc queries is the procedure cache. The query in Listing 5-4 will give you information about the number of execution plans in the procedure cache.

Listing 5-4. Query procedure cache

```
SELECT
    objtype,
    COUNT_BIG (*) AS 'Total Plans',
    SUM(CAST(size_in_bytes AS DECIMAL(12,2)))/1024/1024 AS 'Size (MB)'
FROM sys.dm_exec_cached_plans
GROUP BY objtype;
```

The results of this query should look like Figure 5-13, though the numbers will be different on your system.

	objtype	Total Plans	Size (MB)
1	UsrTab	1	0.218750
2	Prepared	7904	3209.601562
3	View	171	15.906250
4	Adhoc	58308	500.850509
5	Check	2	0.031250
6	Proc	67	21.015625

Figure 5-13. Results of querying procedure cache

We should focus on the number of ad-hoc execution plans. If you see this number growing rapidly and experience CMEMTHREAD waits, it might be worth the effort to analyze some of those ad-hoc queries. If possible, try to optimize the queries so they generate a reusable plan. If your application uses many dynamic queries then try to use the `sp_executesql` system-stored procedure instead of the EXECUTE (EXEC) command. Using the EXEC command will most likely result in a plan that will only be used once.

Microsoft has released various fixes (most notably the partitioning of certain memory objects across CPUs) for this problem in SQL Server 2005 SP2, making it less common these days. Even if you are using newer SQL Server editions than SQL Server 2005, it might be a good idea to upgrade to the latest available Service Pack since there have been various memory-related bug fixes in every SQL Server edition.

CMEMTHREAD Summary

The CMEMTHREAD Wait Type is a memory-related Wait Type. CMEMTHREAD waits occur when multiple threads try to access memory objects that can only be accessed by one thread at a time. The time other threads spend waiting for their turn to access the memory object is recorded as CMEMTHREAD Wait Time. One of the more common cases where CMEMTHREAD

waits can occur is when your system uses a high amount of ad-hoc queries. Every time a new execution plan is generated, SQL Server will access a memory object; if many execution plans are generated this can lead to a queue of threads that want access to the memory object, resulting in CMEMTHREAD waits.

IO_COMPLETION

Just like the ASYNC_IO_COMPLETION Wait Type, IO_COMPLETION waits occur when SQL Server is waiting for storage-related actions to complete. And just like the ASYNC_IO_COMPLETION Wait Type, seeing high Wait Times of the IO_COMPLETION Wait Type doesn't necessarily mean there is something wrong with your storage system. IO_COMPLETION waits occur normally while your SQL Server instance is running and should only be a concern if Wait Times are a lot higher than normal.

What Is the IO_COMPLETION Wait Type?

While the ASYNC_IO_COMPLETION Wait Type is recorded when database-related actions are performed, like a database backup, IO_COMPLETION waits occur when non-data pages are involved, like the restore of a transaction log backup or when bitmap allocation pages, like the GAM page, are accessed. IO_COMPLETION waits can also occur when queries are being executed that perform read or write operations to the storage subsystem, like a Merge Join operator.

IO_COMPLETION Example

Let's generate some IO_COMPLETION waits by restoring a transaction log backup. For this example we will make use of the AdventureWorks database again. The query in Listing 5-5 will perform a full backup of the AdventureWorks database, make some changes, and perform a transaction log backup. When that is complete we will restore the full backup again, clear the `sys.dm_os_wait_stats` DMV, restore the transaction log backup, and check for IO_COMPLETION waits.

Listing 5-5. Generate IO_Completion waits

```
-- Make sure AdventureWorks is in Full recovery model
ALTER DATABASE AdventureWorks SET RECOVERY FULL
GO

-- Perform full backup
BACKUP DATABASE [AdventureWorks]
TO DISK = N'F:\Backup\AW_Full.bak'
GO

-- Make some changes to AW database
USE AdventureWorks
GO
```

```

UPDATE Person.Address
SET City = 'Portland'
WHERE City = 'Bothell'

-- Backup Transaction Log
BACKUP LOG [AdventureWorks]
TO DISK = N'F:\Backup\AW_Log.trn'
GO

-- Restore the previous full backup with NORECOVERY
USE [master]
GO

RESTORE DATABASE [AdventureWorks]
FROM DISK = N'F:\Backup\AW_Full.bak'
WITH NORECOVERY, REPLACE
GO

-- Clear sys.dm_os_wait_stats
dbcc sqlperf ('sys.dm_os_wait_stats', CLEAR)

-- Restore last Transaction Log backup
RESTORE LOG [AdventureWorks] FROM DISK = N'F:\Backup\AW_Log.trn'
GO

-- Check IO_COMPLETION waits
SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'IO_COMPLETION'

```

The results of this query against the `sys.dm_os_wait_stats` DMV on my test system can be seen in Figure 5-14.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	IO_COMPLETION	266	188	16	0

Figure 5-14. *IO_COMPLETION* waits

We only modified a few records using the query in Listing 5-5, so the total Wait Time is pretty low.

`IO_COMPLETION` Wait Times also occur when you are starting up your databases after, for instance, a restart of the SQL Server service. This means you should expect `IO_COMPLETION` waits after a restart or a failover; these are completely normal.

Lowering IO_COMPLETION Waits

Most of the time IO_COMPLETION waits shouldn't be a cause for concern. When they are a lot higher than the Wait Times in your baseline, you should analyze the storage subsystem performance like I described in the “Lowering ASYNC_IO_COMPLETION” section. While certain query operations can also cause IO_COMPLETION waits, these are frequently not the cause for higher-than-normal Wait Times.

IO_COMPLETION Summary

Just like the ASYNC_IO_COMPLETION Wait Type, IO_COMPLETION waits occur when accessing your storage subsystem. IO_COMPLETION waits occur when SQL Server is waiting on non-data page operations to complete, like a transaction log restore operation or the reading of bitmap pages, like the GAM page. Seeing waits of the IO_COMPLETION type occur is completely normal and these frequently do not require deeper analysis unless the Wait Times are a lot higher than the values in your baseline. In those cases focus on the performance (and especially latency) of your storage subsystem first.

LOGBUFFER & WRITELOG

I have combined both LOGBUFFER and WRITELOG in this section. This is because both Wait Types have a close relation to each other. Both of them are related to the transaction log and the storage subsystem.

What Are the LOGBUFFER & WRITELOG Wait Types?

To understand what the LOGBUFFER and WRITELOG Wait Types represent, we need to have some understanding of how SQL Server writes to the transaction log. In short, the following events happen whenever we change or add data inside a database:

1. Data page where the data resides is modified in the buffer cache; if the page wasn't already in the buffer cache, it will get read into the buffer cache first.
2. The data page will be marked as “dirty” inside the buffer cache.
3. The log records that represent the modification get saved in the log buffer.
4. A log flush occurs (this can be for multiple reasons, which we will discuss later), writing the log records from the log buffer to the transaction log.
5. The dirty data page gets written to the data file.

To show this behavior I have included Figure 5-15.




Figure 5-15. How a transaction moves

The action representing the movement of a dirty data page, and the action of writing the log records to the transaction log, are shown as dashed lines. I did this on purpose to illustrate that both of these actions do not necessarily happen directly.

As you probably know, dirty data pages are updated inside the buffer cache first, and are only written to the data file when a checkpoint operation occurs. This means that a dirty page can stay inside the memory even after your transaction was committed.

This is not true for log records inside the log buffer. As soon as your transaction commits, and that transaction has an active log record in the log buffer, all the log records inside that log buffer are written (or flushed) to the transaction log on disk. But this doesn't only occur when the transaction is committed. The log buffer has a fixed size of 60 KB, and as soon as the log buffer is full, it will flush all the records inside it to the transaction log.

Let's add both the Wait Types we are discussing in this section to the story. The `WRITELOG` Wait Type occurs whenever SQL Server is flushing the contents of the log buffer to the transaction log on disk. The `LOGBUFFER` Wait Type occurs when inserting log records in the log buffer, when at the time of insertion SQL Server has to wait for free space inside the log buffer. I have added both the Wait Types at the parts where they can get generated in Figure 5-16.

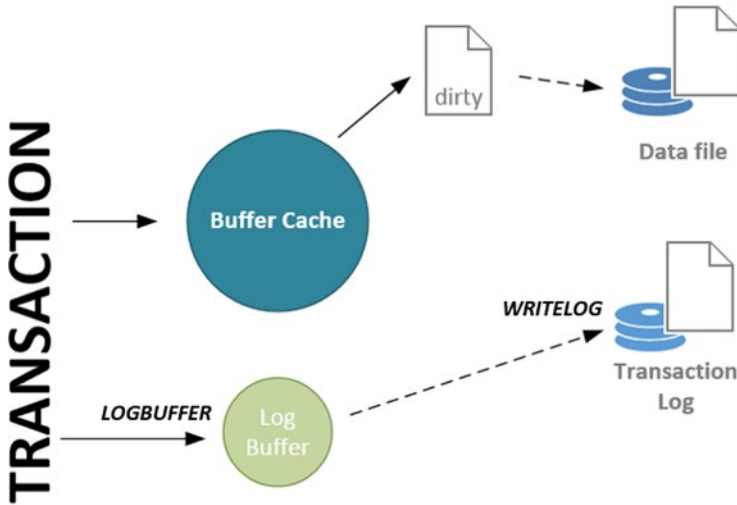


Figure 5-16. Transaction movement and the LOGBUFFER & WRITELOG Wait Types

You can now probably see how both Wait Types are related. Whenever a long WRITELOG wait occurs, chances are you will also see LOGBUFFER waits if the process that is writing the log records to the transaction log on disk cannot process them as fast as the log records enter the log buffer.

This situation frequently occurs on systems with a lot of concurrent data modifications. This results in a high volume of transactions that need to be written to disk. Another common cause is the performance of the storage subsystem where the transaction log file resides. If the storage subsystem has sub-optimal performance, your WRITELOG Wait Times will increase, with the possibility existing that LOGBUFFER waits can occur if the volume of transactions is high enough.

The performance of your transaction log is critical for the performance of your entire database. Slow transaction log performance will have an impact on every change you perform inside your database, as every modification has to be written to the transaction log before it can get committed.

LOGBUFFER & WRITELOG Example

To give you an example of LOGBUFFER and WRITELOG waits occurring, I am creating a new database using the script in Listing 5-6.

Listing 5-6. Create trans_demo database

```

USE master
GO

-- Create demo database
CREATE DATABASE [trans_demo]
ON PRIMARY
(
    NAME = N'trans_demo', FILENAME = N'D:\Data\trans_demo.mdf' ,
    SIZE = 153600KB , FILEGROWTH = 10%
)
LOG ON
(
    NAME = N'trans_demo_log', FILENAME = N'D:\Log\trans_demo.ldf' ,
    SIZE = 51200KB , FILEGROWTH = 10%
)
GO

-- Make sure recovery model is set to full
ALTER DATABASE [trans_demo] SET RECOVERY FULL
GO

-- Create a simple test table
USE trans_demo
GO

CREATE TABLE transactions
(
    t_guid VARCHAR(50)
)
GO

```

Now that we have created a brand new database, I am going to use the Ostress utility to generate load against the trans_demo database. I am going to execute the query in Listing 5-7, which I saved to the logbuffer_impl.sql file, with 200 concurrent connections using the following command "C:\Program Files\Microsoft Corporation\RMLUtils\ostress.exe" -E -dtrans_demo -i"C:\logbuffer_impl.sql" -n200 -r1 -q.

Listing 5-7. Insert rows inside the trans_demo database

```

DECLARE @i INT
SET @i = 1

WHILE @i < 10000

    BEGIN

```

```

INSERT INTO transactions
    (t_guid)
VALUES
    (newid())

SET @i = @i + 1

END

```

Before I started the Ostress utility, I cleared the `sys.dm_os_wait_stats` DMV. After about three minutes on my test SQL Server, the Ostress utility finished executing the workload. If I query the `sys.dm_os_wait_stats` DMV and look for the `LOGBUFFER` and `WRITELOG` Wait Types, I get the results shown in Figure 5-17.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms
1	WRITELOG	2091319	17520781	4937
2	LOGBUFFER	760	255438	4930

Figure 5-17. *WRITELOG & LOGBUFFER waits*

Lowering LOGBUFFER & WRITELOG Waits

There are generally two approaches you can take to lowering `LOGBUFFER` and `WRITELOG` waits, keeping in mind though that `WRITELOG` waits do also occur normally and its Wait Time should only be a cause for concern if it is a lot higher than normal.

The first approach is to take a good look at how your transactions are being executed. In the example above we implicitly committed every `INSERT` statement. This means that as soon as the log record for the `INSERT` statement entered the log buffer it needed to be flushed again. If we would explicitly commit the whole `WHILE` loop we would have larger writes to flush to the transaction log, resulting in better performance. This is because writing small blocks frequently is generally slower than writing large blocks at a larger interval. Cursors can also have the same effect as the example we used, so use them as little as possible.

The other approach is based on the storage subsystem. If SQL Server cannot write the log records fast enough, you can encounter `LOGBUFFER` and `WRITELOG` waits. As a best practice, make sure to split your transaction log and database data files on separate disks so they do not impact each other in times of heavy load. Also monitor the disk the transaction log is located on using the disk-performance counters in Perfmon, like `Avg. Disk sec/write`, to show you the write latency, and `Disk Writes/sec`, to show the write IOPS and check if the values are inside the acceptable range.

If your SQL Server instance is running SQL Server 2014, you could choose to make use of the `Delayed Durability` option, which was introduced in SQL Server 2014. In short, enabling this option will no longer flush the log buffer content to disk when a transaction commits, but rather will wait until the log buffer is full (60 KB) before flushing the contents to the transaction log. By enabling this option you are running a risk that transactions that have committed, but have not yet been written to the transaction log, will be lost during a failure since they will only be written to the transaction log when the log buffer is full.

LOGBUFFER & WRITELOG Summary

Both the LOGBUFFER and WRITELOG Wait Types are related to the way SQL Server processes transactions. The WRITELOG Wait Type occurs every time a log records is written to the transaction log, and generally isn't a cause for concern. When paired with high LOGBUFFER Wait Times, high WRITELOG Wait Times can indicate transaction log pressure. To lower those Wait Times try to avoid cursor and WHILE statements, since the statements inside the cursor or WHILE clauses will often get implicitly committed, creating a large amount of small writes. Also check your storage configuration to make sure the transaction log isn't on the same drive as the database data file. If high Wait Times still occur, analyze the performance of the disk on which your transaction log is located.

RESOURCE_SEMAPHORE

The RESOURCE_SEMAPHORE Wait Type is a memory-related Wait Type that can show itself when a query-memory request cannot be granted immediately. These waits can occur on servers that are experiencing memory pressure, or when a great number of concurrent queries request memory for expensive operations like sorts or joins.

What Is the RESOURCE_SEMAPHORE Wait Type?

When a query is executed in SQL Server a series of steps occur before the actual execution. The first step involved is that a compiled plan is generated. This plan contains the logical instructions, or operations, needed to fulfill the query requests. During the generation of the compiled plan, a calculation is performed to determine the amount of memory needed to execute the query, which depends on the operations involved in the compiled plan. Some of the operations that require memory are sorts and joins, which temporarily store row data in the memory of the SQL Server. The minimum amount of memory needed to perform these sorts or joins is known as the *required memory*, without which the query simply cannot get executed. If more memory is needed to store row data in memory during a sort, for instance, it will be calculated as *additional memory*. Without this additional memory, a query can still get executed, but instead of writing the temporary row data to memory it will write it to disk.

When the query gets executed a memory grant will be determined based on the required and additional memory values calculated in the compiled plan. This memory grant is needed in order to perform a memory reservation at an internal object called the *resource semaphore*. The resource semaphore is responsible for reserving the memory a query needs for execution, but it also manages memory throttling when too many queries concurrently ask for memory reservations or when there is not enough memory available at that time. It does this by maintaining a queue of queries that are requesting memory. If there are no queries inside the queue and a new query requests memory, the resource semaphore will grant it to the query (if enough free memory is available). However, if there is a queue the new query will be put at the end of the queue, and it has to wait for its turn to receive a memory grant.

Before the resource semaphore will grant the requested memory to a query it will check whether there is enough free memory to execute it. If, for some reason, there is less memory available than the amount requested by the query, the query will get put in the queue again until enough memory is available. When a query is inside the resource semaphore queue waiting for its requested memory, the time it spends inside the queue will be recorded as the RESOURCE_SEMAPHORE Wait Type.

There is a maximum amount of memory available for the resource semaphore to use, and it is allocated from the buffer cache. The resource semaphore can allocate up to 75% of the memory from the buffer cache for memory grants, but a single query can never get more than 25% of that amount. For instance, if we have a SQL Server with a buffer cache that can grow to 500 MB, we would have a maximum of 375 MB for memory grants. A single query in this example can never receive more than 93 MB. Having that much memory being possibly granted to queries can be problematic, since that memory is not being used for the buffer cache, meaning more IOs to the storage subsystem are needed to retrieve and write data pages.

RESOURCE_SEMAPHORE Example

For this example we are going to execute a query against the AdventureWorks database that involves a sort operation. As I mentioned in the previous section, a query that involves a sort will request memory from the resource semaphore so as to perform the sort operation. We will also use the Ostress tool to create a situation where multiple queries are requesting memory, creating a queue at the resource semaphore.

Let's take a look at the query and the memory grant information that we will be executing in Listing 5-8.

Listing 5-8. Sort query against the AdventureWorks database

```
SELECT
    SalesOrderID,
    SalesOrderDetailID,
    ProductID,
    CarrierTrackingNumber
FROM
    Sales.SalesOrderDetail
ORDER BY CarrierTrackingNumber ASC
```

As you can see, this is a relatively simple query that returns some information from the Sales.SalesOrderDetail table, ordered by the CarrierTrackingNumber.

If we enable the Include Actual Execution Plan option and execute the query, we can take a look at the amount of memory that was needed to execute it. The results on my test SQL Server are shown in Figure 5-18. You can access these properties by showing the Properties window (View - Properties Window) or by pressing F4, and selecting the SELECT operator.

MemoryGrantInfo	
DesiredMemory	13568
GrantedMemory	13568
GrantWaitTime	0
MaxUsedMemory	8608
RequestedMemory	13568
RequiredMemory	512
SerialDesiredMemory	13568
SerialRequiredMemory	512

Figure 5-18. *MemoryGrantInfo* inside the properties of the execution plan

Because we requested the actual execution plan we can also see the amount of memory that was granted to the query for execution. In this case the query got 13,568 KB (13.25 MB) granted by the resource semaphore, as shown in the `GrantedMemory` property. The minimal amount of memory needed to execute the query, the required memory, was 512 KB, shown by the `RequiredMemory` property. The query asked for 13,568 KB, as shown in the `DesiredMemory` property, which is the sum of the required and additional memory. We can see the query received what it asked for, since both the `GrantedMemory` and `DesiredMemory` have the same value.

There are two other properties in Figure 5-18 I would like to point out, the `SerialDesiredMemory` and the `SerialRequiredMemory` properties. In the case of this query, both these properties have the same values as the `DesiredMemory` and `RequiredMemory` properties. This is because the query was performed without using parallelism. When you use parallelism in your queries, more memory is needed to perform the sort operation since work is split up among threads. Figure 5-19 shows the `MemoryGrantInfo` properties when I forced the query in Listing 5-8 to use parallelism, spreading the work among four threads.

MemoryGrantInfo	
DesiredMemory	15488
GrantedMemory	15488
GrantWaitTime	0
MaxUsedMemory	9024
RequestedMemory	15488
RequiredMemory	2432
SerialDesiredMemory	13568
SerialRequiredMemory	512

Figure 5-19. *MemoryGrantInfo* properties when executing a parallel query

As you can see, the `SerialRequiredMemory` has the same value as when we executed the query serially. The `RequiredMemory` and `RequestedMemory` have increased in size so that the sort operation can be completed using parallelism. You should keep this information in mind when you run into memory-related issues and when many of your queries involve sort and join operations that are performed using parallelism, since parallelism simply requires more memory.

Now that we know how much memory is needed to execute the query in Listing 5-8, let's use Ostress to execute the query using multiple connections. Before I start Ostress, I must change the maximum server memory value to 250 MB using the following query:

```
EXEC sys.sp_configure N'max server memory (MB)', N'250'
GO
```

```
RECONFIGURE WITH OVERRIDE
GO
```

I saved the query in Listing 5-8 to a .sql file named `resource_semaphore.sql` and executed Ostress using the following command line: "C:\Program Files\Microsoft Corporation\RMLUtils\ostress.exe" -E -dAdventureWorks -i"C:\resource_semaphore.sql" -n20 -r1 -q

This will execute the `resource_semaphore.sql` script against the AdventureWorks database with 20 concurrent connections, with each connection performing the query one time.

While Ostress is running, I query the `sys.dm_os_waiting_tasks` DMV, looking for RESOURCE_SEMAPHORE waits; some of the results are shown in Figure 5-20.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address
4	0x000000017F013C38	62	0	594	RESOURCE_SEMAPHORE	0x000007FEF5D15B00
5	0x000000017F013C38	62	0	594	RESOURCE_SEMAPHORE	0x000007FEF5D15B00
6	0x000000017F013C38	62	0	594	RESOURCE_SEMAPHORE	0x000007FEF5D15B00
7	0x0000000174352558	61	0	587	RESOURCE_SEMAPHORE	0x000007FEF5D15B00
8	0x0000000174352558	61	0	587	RESOURCE_SEMAPHORE	0x000007FEF5D15B00
9	0x0000000174352558	61	0	587	RESOURCE_SEMAPHORE	0x000007FEF5D15B00
10	0x0000000174352558	61	0	587	RESOURCE_SEMAPHORE	0x000007FEF5D15B00

Figure 5-20. RESOURCE_SEMAPHORE waits in the `sys.dm_os_waiting_tasks` DMV

Because we set our maximum server memory to 250 MB, and each query requests 13.25 MB memory, we do not have enough memory free to grant all the memory requested. This will result in the RESOURCE_SEMAPHORE Wait Type you can see in Figure 5-20.

There are various other resources we can use to analyze RESOURCE_SEMAPHORE waits. The resource semaphores themselves have their own DMV, `sys.dm_exec_query_resource_semaphores`, which will return information about their memory consumption and outstanding and waiting grants. Figure 5-21 shows the results of the query that follows against the `sys.dm_exec_query_resource_semaphores` DMV, while running the Ostress workload:

```
SELECT
    target_memory_kb,
    max_target_memory_kb,
    total_memory_kb,
    available_memory_kb,
    granted_memory_kb,
```

```

grantee_count,
waiter_count
FROM sys.dm_exec_query_resource_semaphores
WHERE pool_id = 2

```

	target_memory_kb	max_target_memory_kb	total_memory_kb	available_memory_kb	granted_memory_kb	grantee_count	waiter_count
1	109216	189256	122112	0	122112	9	11
2	5744	NULL	5744	5744	0	0	0

Figure 5-21. *sys.dm_exec_query_resource_semaphores*

I am filtering out `pool_id` 1, because this pool will not handle user queries.

As you might have noticed, two rows are returned. This is because there are actually two different resource semaphores. The top row is the “regular” resource semaphore. This will handle queries that request more than 5 MB memory. The second row (identified by the NULL value of the `max_target_memory_kb` column) returns information for the “small” resource semaphore, which handles queries that are smaller than 5 MB. Because our query requested more than 5 MB of memory, we will receive our memory grants from the regular resource semaphore.

Let’s go through the various columns that are returned by the query against the `sys.dm_exec_query_resource_semaphore` DMV:

- The `target_memory_kb` column returns the amount of memory in KB that this resource semaphore plans to use as a maximum amount of memory it can grant to queries.
- The `max_target_memory_kb` column returns the maximum amount of memory this resource semaphore could grant.
- The `total_memory_kb` column returns the total memory held by the resource semaphore and is the sum of the `available_memory_kb` and the `granted_memory_kb`.
- The `granted_memory_kb` returns the amount of memory that is granted to queries at this time.
- The `grantee_count` and `waiter_count` columns return the amount of grants that have currently been satisfied or are waiting in the resource semaphore queue.

From this information we can see that the information returned by the `granted_memory_kb` column is correct, and that our test queries are requesting the memory grants. We know from the execution plan that our test query will request 13,568 KB. Since the `grantee_count` column shows us that 9 memory requests are granted, we can multiply the amount of memory requests with the amount of memory per query (9 * 13,568 KB), which ends up being 122,112 KB, the amount of granted memory in Figure 5-21.

We can also use Perfmon to monitor the total size of the granted memory by looking at the SQLServer:Memory Manager\Granted Workspace Memory (KB) counter, as shown in Figure 5-22.

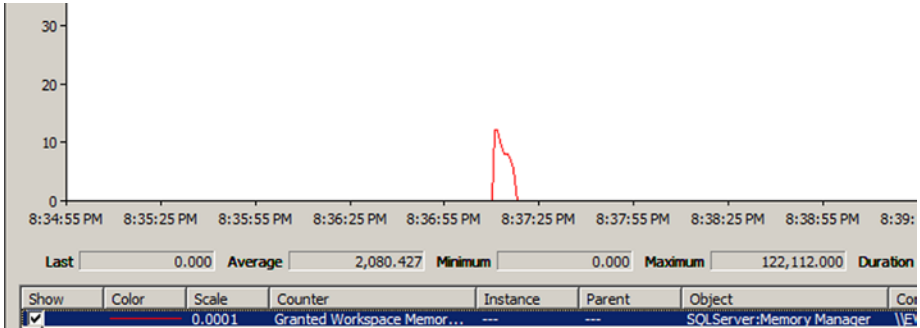


Figure 5-22. Granted Workspace Memory (KB) Perfmon counter

Notice the spike in Figure 5-22, which occurred when I executed the Ostress workload, and just as in the results in Figure 5-21, the amount of memory granted was 122,112 KB.

Lowering RESOURCE_SEMAPHORE Waits

There are various possible methods you can use to lower, or even resolve, the Wait Times of the RESOURCE_SEMAPHORE Wait Type. The most obvious one would be adding more memory, but this can end up being an expensive solution while other less expensive options exist.

The first possible solution would be to look at the queries that are requesting large amounts of memory for their execution. You should focus on the queries that are performing large sorts or joins (especially hash joins) and check whether you can lower the number of rows that need to be sorted or joined, or avoid the sort or join completely. One way to avoid a sort operation would be to add an index to the table where the sort is performed. If the order of values inside the index were the same as the sort operation, a sort operation would no longer be necessary, since the index would already have ordered the results.

Another solution involves parallelism. If queries use parallelism during sort or join operations, more memory is requested than when the query is executed serially. Modifying queries so they don't use parallelism, by either using query hints or changing the parallelism configuration for the whole SQL Server instance, will result in lower amounts of memory being required to execute the queries.

Finally, if you are running an Enterprise Edition of SQL Server, you could use the resource governor feature to configure the memory usage of each resource pool. By configuring the amount of memory a certain resource pool can use, you can also set the amount of memory a resource semaphore can grant. We won't go into detail about the resource governor feature in this book, but more information can be found on the MSDN page of the resource governor with the following hyperlink: <https://msdn.microsoft.com/en-us/library/bb933866.aspx>.

RESOURCE_SEMAPHORE Summary

The RESOURCE_SEMAPHORE Wait Type is related to the amount of memory a query needs to perform certain operations, like sorts and joins. An object, named a resource semaphore, is responsible for managing and throttling the memory requests of queries. If a query requests more memory than the resource semaphore can grant, the memory request will be moved into the resource semaphore queue. While the memory request is inside the resource semaphore queue, RESOURCE_SEMAPHORE Wait Times are recorded. There are various methods to lower or resolve RESOURCE_SEMAPHORE waits. You can choose to add more memory to the SQL Server instance, or optimize the queries so sorts and joins do not require as much memory. Another option is using the resource governor, where you can define resource pools so as to minimize the impact of large memory requests.

RESOURCE_SEMAPHORE_QUERY_COMPILE

In the previous section we discussed the RESOURCE_SEMAPHORE Wait Type, which indicates that there is not enough free memory available for certain query operations like sorts and joins. Just like the RESOURCE_SEMAPHORE Wait Type, the RESOURCE_SEMAPHORE_QUERY_COMPILE Wait Type is also related to the memory of your SQL Server instance. But instead of indicating a shortage in query memory, the RESOURCE_SEMAPHORE_QUERY_COMPILE Wait Type indicates a memory shortage during the compilation process of the query.

What Is the RESOURCE_SEMAPHORE_QUERY_COMPILE Wait Type?

In the explanation of the RESOURCE_SEMAPHORE Wait Type we discussed what resource semaphores are and what they do. For the explanation of the RESOURCE_SEMAPHORE_QUERY_COMPILE Wait Type we are going to dive a little deeper into the inner workings of the resource semaphore.

You should think of resource semaphores as “gateways” that throttle direct access to memory resources. There are different tasks a resource semaphore can do. In the previous section we discussed the resource semaphores that were responsible for granting memory for certain operations, like sorts and joins. We also noted that there are two resource semaphores that are responsible for granting this memory—the regular resource semaphore that handles queries that request 5 MB or more memory, and the small resource semaphore that handles memory grants for queries that request less than 5 MB memory.

The resource semaphores that are related to the RESOURCE_SEMAPHORE_QUERY_COMPILE Wait Type are responsible for memory grants that are needed during the compilation process of a query, excluding the memory needed for query execution. Just like the resource semaphores in the previous section, the ones responsible for memory grants during the compilation process also have different gateways. Figure 5-23 shows the different gateways for the compilation-memory resource semaphore.

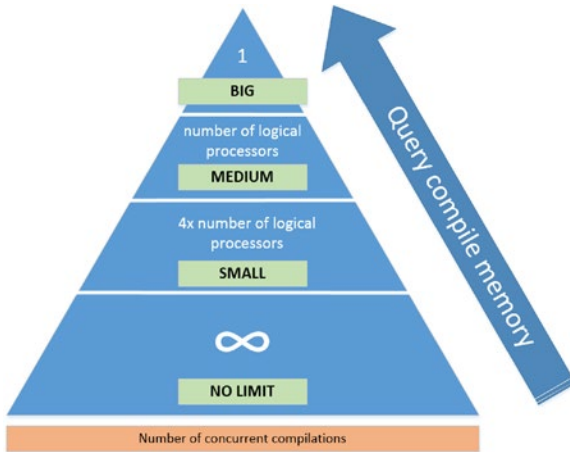


Figure 5-23. *Compilation-memory resource semaphore*

By default there are three gateways: small, medium, and big. Depending on the amount of memory the compilation of a query needs, it will get assigned to either of the three. If the amount of memory needed for compilation is less than the memory threshold for the small gateway, the query does not have to pass through a gateway. The amount of concurrent compilations, or queries, that can move through the gateway simultaneously is calculated by the number of logical processors available for your SQL Server instance. For example, if your SQL Server instance has 4 logical processors, the small gateway will allow 16 concurrent compilations and the medium gateway 4. The big gateway will always only allow 1 query at a time to compile.

The memory threshold for the small gateway is static, but for the medium and big gateways the thresholds are dynamic. This means that the compilation memory needed to reach the medium or big gateways can change during the runtime of your SQL Server instance.

The whole purpose of these gateways is to ensure that the need for compilation memory stays under control. This avoids out-of-memory situations in cases where many large compilation-memory requests would automatically be granted and would drain the SQL Server instance of its memory.

Before we continue and take a look how we can access gateway information from inside SQL Server, let's go through an example of a query compilation.

Say we have a query that needs 1560 KB of compilation memory. The query will start by requesting a gateway. Since our small gateway has a threshold of 370 KB and our medium gateway has a threshold of 5346 KB, the query will end up in the small gateway. If there are any queries currently in a queue at the small gateway, the query will enter the queue and wait until its turn, all the time logging `RESOURCE_SEMAPHORE_QUERY_COMPILE` Wait Time. While the query is getting compiled, the amount of memory used during the compilation is tracked; if the query ends up using more memory and reaches the threshold of the medium gateway, it will get moved to the medium gateway. When the query compilation is finished, it will be removed from the gateway.

We can access information about the resource semaphore gateways from inside SQL Server by executing the DBCC MEMORYSTATUS command. Somewhere in the enormous amount of results you will find the gateway information as shown in Figure 5-24.

	Small Gateway (default)	Value
1	Configured Units	16
2	Available Units	16
3	Acquires	0
4	Writers	0
5	Threshold Factor	380000
6	Threshold	380000

	Medium Gateway (default)	Value
1	Configured Units	4
2	Available Units	4
3	Acquires	0
4	Writers	0
5	Threshold Factor	12
6	Threshold	-1

Figure 5-24. Gateway information returned by the DBCC MEMORYSTATUS command

Let's go through the results that are returned for the gateways.

The Configured Units row returns the maximum amount of concurrent compilations allowed for this gateway. This is determined by the number of logical processors available for your SQL Server instance. Because my test SQL Server has 4 logical processors, I have 16 slots for the small gateway (4 * logical number of processors), and 4 for the medium gateway. The Available Units row shows the number of currently free slots for this gateway, while the Acquires row shows the slots currently taken by compilations. The amount of queries that have to wait for a free slot are shown in the Writers row. The Threshold value is the amount of memory in bytes that a query compilation would need in order to enter the gateway. For my test SQL Server system, the small gateway has a threshold of 380,000 bytes, or 371 KB. As you might notice in Figure 5-24, the medium gateway has a threshold of -1. This is because of the dynamic nature of the thresholds of the medium and big gateways. Since there is no activity at the gateway below the medium one, there is no need to set a threshold yet.

RESOURCE_SEMAPHORE_QUERY_COMPILE Example

To show you an example of RESOURCE_SEMAPHORE_QUERY_COMPILE waits in action, I am going to execute the query in Listing 5-9 multiple times, using many concurrent connections. The query is a dynamic query that selects a random row from

two joined tables inside the AdventureWorks database. In this case it doesn't matter if any results are returned or not—the thing we are trying to achieve here is the creation of compilation-memory contention.

Listing 5-9. RESOURCE_SEMAPHORE_QUERY_COMPILE wait query

```

DECLARE @ID VARCHAR(250)
DECLARE @SQL VarChar(MAX)
SET @ID = FLOOR(RAND()*(20000-1)+1);

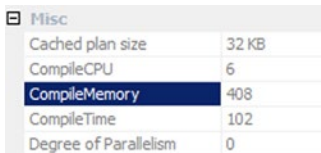
SET @SQL =
    '
    SELECT
        ' + @ID + ',
        SUM(soh.SubTotal),
        COUNT(soh.SubTotal)
    FROM sales.SalesOrderHeader soh
    INNER JOIN person.Person p
        ON soh.SalesPersonID = p.BusinessEntityID
    WHERE p.BusinessEntityID = ' + @ID + '
    '

EXEC (@SQL)

```

Before we execute the query with many concurrent connections, let's check how much compilation memory would be needed. We can do this by executing the query in Listing 5-9 in SQL Server Management Studio and enabling the actual execution plan.

After executing the query and opening the actual execution plan, we need to look at the CompileMemory property. You can access these properties by showing the Properties window (View - Properties Window) or by pressing F4, and selecting the SELECT operator. Figure 5-25 shows the actual execution plan properties on my Test SQL Server.



Misc	
Cached plan size	32 KB
CompileCPU	6
CompileMemory	408
CompileTime	102
Degree of Parallelism	0

Figure 5-25. Miscellaneous execution plan properties

The value returned by the CompileMemory property is the amount of compile memory needed expressed in KB. For this query 408 KB is needed for compilation. The threshold for the small gateway on my Test SQL Server was 371 KB, so I am pretty sure the query will access the small gateway.

Again, we are going to use the Ostress utility to generate the needed concurrent connections to execute the query. I saved the query to the resource_semaphore_compile.sql file, and then used that file as input for the Ostress command below. Because the query is very fast, I let every connection execute it 100 times so that we have some time to look at the Wait Statistics.

```
"C:\Program Files\Microsoft Corporation\RMLUtils\ostress.exe" -E
-dAdventureWorks -i"C:\resource_semaphore_compile.sql" -n200 -r100 -q
```

After a few seconds many (around 1700) RESOURCE_SEMAPHORE_QUERY_COMPILE waits can be seen in the sys.dm_os_waiting_tasks DMV, as shown in Figure 5-26.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type
79	0x000000016CD02C...	182	0	30	RESOURCE_SEMAPHORE_QUERY_COMPILE
80	0x000000016CD02C...	182	0	30	RESOURCE_SEMAPHORE_QUERY_COMPILE
81	0x000000016B176558	158	0	241	RESOURCE_SEMAPHORE_QUERY_COMPILE
82	0x000000016B176558	158	0	241	RESOURCE_SEMAPHORE_QUERY_COMPILE
83	0x000000016B176558	158	0	241	RESOURCE_SEMAPHORE_QUERY_COMPILE
84	0x000000016B176558	158	0	292	RESOURCE_SEMAPHORE_QUERY_COMPILE

Figure 5-26. RESOURCE_SEMAPHORE_QUERY_COMPILE waits

If we now execute the DBCC MEMORYSTATUS command we should be able to find out at what gateway the compilation contention is occurring. Figure 5-27 shows the gateway output of the DBCC MEMORYSTATUS command on my test SQL Server.

	Small Gateway (default)	Value
1	Configured Units	16
2	Available Units	0
3	Acquires	15
4	Waiters	3
5	Threshold Factor	380000
6	Threshold	380000

	Medium Gateway (default)	Value
1	Configured Units	4
2	Available Units	4
3	Acquires	0
4	Waiters	0
5	Threshold Factor	12
6	Threshold	537577

Figure 5-27. DBCC MEMORYSTATUS during compilation contention

As you can see in Figure 5-27, if we look at the amount of Available Units, there are no available slots left for new compilation-memory requests. As a matter of fact, we have three compilation-memory requests waiting in the resource semaphore queue. Also note that the threshold of the medium gateway has now changed from -1 to 537577 bytes (525 KB). Now that contention is occurring on a lower gateway, the threshold for the medium gateway is dynamically determined, even though there are no compilation-memory requests being processed by this gateway.

Lowering RESOURCE_SEMAPHORE_QUERY_COMPILE Waits

The methods you can use to lower the Wait Times of the RESOURCE_SEMAPHORE_QUERY_COMPILE Wait Type are in many cases the same as those that you would use to lower or resolve RESOURCE_SEMAPHORE waits. Just like the RESOURCE_SEMAPHORE Wait Type, the RESOURCE_SEMAPHORE_QUERY_COMPILE Wait Type is memory related, so if you can increase the total amount of memory available for query compilation, chances are you will lower or resolve RESOURCE_SEMAPHORE_QUERY_COMPILE Wait Times. Increasing memory is, however, in many cases the last resort.

Because we can access very specific information about the gateways of the resource semaphore that is dealing with the compilation memory by using the DBCC MEMORYSTATUS command, a good first step is to analyze the usage patterns of the gateways. If you notice that one specific gateway constantly has waiting memory requests, then the memory threshold of that gateway, or the maximum allowed amount of concurrent compilation-memory requests, should give you some hints about the root cause. For instance, if you notice many queued compilation-memory requests at the big gateway (which only allows one query at a time), the source of your RESOURCE_SEMAPHORE_QUERY_COMPILE Wait Times may be the queries that request a large amount of compilation memory. Another cause may be a large number of concurrent queries that all need to access the small gateway, which was the case in our example, causing a queue at the gateway.

In these cases you should find the specific queries that cause the queues at the gateways and try to optimize them, either by lowering the amount of compilation memory or by making sure fewer compilations happen. The latter can be done by making sure your queries are being parameterized correctly. Queries that generate ad-hoc plans every time they are executed can be a cause of RESOURCE_SEMAPHORE_QUERY_COMPILE waits, especially if they are executed very frequently and concurrently. Jonathan Kehayias of SQLskills has written an excellent blog post on how you can query the plan cache to detect heavy compilation queries; it can be found over at <https://www.sqlskills.com/blogs/jonathan/identifying-high-compile-time-statements-from-the-plan-cache/>. Using the script in Jonathan's blog post should help you with detecting those queries that require a large amount of compilation memory.

If your SQL Server is under memory pressure, it is also possible to see RESOURCE_SEMAPHORE_QUERY_COMPILE waits occur. This happens because of the dynamic compilation-memory thresholds of the medium and big gateways. If SQL Server is under memory pressure, the thresholds of both these gateways will lower, giving more queries the chance to use the medium or big gateways. But because the medium and big gateways allow fewer concurrent compilations, in the small gateway the available concurrent slots will be filled faster.

Just as with the `RESOURCE_SEMAPHORE` Wait Type, you can use the resource governor to split workloads into specific resource pools. Each resource pool will have its own resource semaphores responsible for granting compilation memory, making it possible to split heavy compilation-memory usage across multiple resource pools.

RESOURCE_SEMAPHORE_QUERY_COMPILE Summary

Just like the resource semaphores that are needed to grant memory requests for specific query operations, resource semaphores exist for access to compilation memory. These resource semaphores throttle access to compilation memory through the usage of gateways. When a query is compiled it will approach a gateway based on the amount of compilation memory it needs. The gateway can then grant the compilation memory requested or put the request in a queue if there are more requests that concurrently want to access the gateway. When a query is waiting inside one of these queues, the `RESOURCE_SEMAPHORE_QUERY_COMPILE` Wait Type is recorded.

Resolving or lowering `RESOURCE_SEMAPHORE_QUERY_COMPILE` Wait Times is commonly achieved by either freeing up more memory or by lowering the compilation-memory needs of queries.

SLEEP_BPOOL_FLUSH

The `SLEEP_BPOOL_FLUSH` Wait Type is directly related to the checkpoint process inside SQL Server. The checkpoint process is responsible for writing modified, or “dirty,” data pages from the buffer pool to the database data file on disk. So next to having a close relationship with the checkpoint process, `SLEEP_BPOOL_FLUSH` waits also have a relationship with the performance of your storage subsystem. If we search for the definition of the `SLEEP_BPOOL_FLUSH` Wait Type on Books Online, Microsoft describes the Wait Type as occurring “when a checkpoint is throttling the issuance of new I/Os in order to avoid flooding the disk subsystem.”

It is pretty common to see `SLEEP_BPOOL_FLUSH` waits occur, and frequently they will not indicate a problem. There are, however, cases where `SLEEP_BPOOL_FLUSH` waits can indicate performance problems that are related to either the checkpoint process or the storage subsystem.

What Is the SLEEP_BPOOL_FLUSH Wait Type?

To get a better understanding of how the `SLEEP_BPOOL_FLUSH` Wait Type gets recorded, we need an understanding of how the checkpoint process works inside SQL Server.

The checkpoint process is an internal SQL Server process that is responsible for writing modified, or dirty, pages from the buffer cache to the database data file. One of the main reasons for this is to speed up recovery of your database when an unexpected failure occurs. When an unexpected failure occurs, SQL Server needs to go back to the state that existed before the failure. It will do this by using the contents of the transaction log to redo, or undo, changes that were made to data pages. If the data page was modified,

but the change was not yet written to the database data file, SQL Server will need to redo the change to the data page. If a checkpoint already wrote the changed data page to the database data file, this step is not needed, which speeds up the recovery process for the database because SQL Server knows the data was written to the database data file. Figure 5-28 shows the (simplified) process that happens when a data page gets modified.

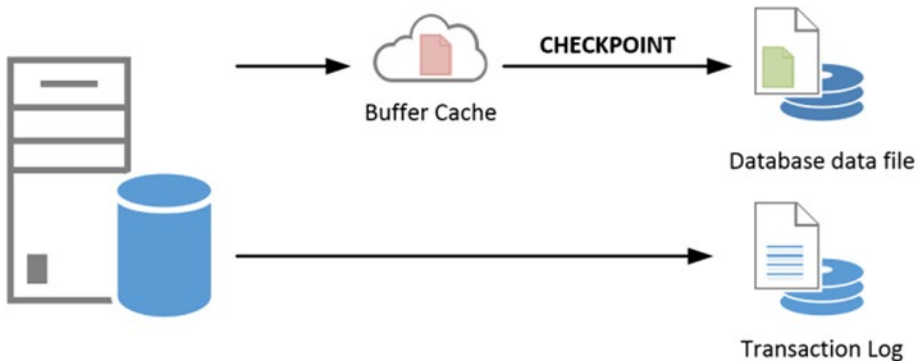


Figure 5-28. Data modification process

The first thing that happens when a data page is modified by a committed transaction is that the change will be recorded in the transaction log (first in the log buffer then to disk as described in the `WRITELOG` & `LOGBUFFER` Wait Types section). The modification of the data page will happen in the buffer cache, and the data page will be marked as dirty (red page icon). When a checkpoint occurs, which can be for multiple reasons as we will discuss later, all of the data pages that have been marked as dirty since the previous checkpoint will be written to the physical database data file on your storage subsystem, regardless of the state of the transaction that created those dirty pages (green page icon).

The checkpoint process is executed by SQL Server automatically roughly once every minute, which is the default recovery time interval. This does not mean that a checkpoint will occur every minute exactly. The values you can specify for the recovery interval are the upper time limit at which a checkpoint should occur, the checkpoint process analyses of the outstanding I/O requests, and latency; throttle checkpoint operations to avoid overloading the storage subsystem.

The list below will describe the various checkpoint types available in SQL Server.

- The internal checkpoint type is not configurable and occurs automatically when certain actions are performed; for instance, a database backup.
- Automatic. These are the default checkpoints that occur roughly every minute when left at their default value of 0. We can change the interval of the checkpoint process by changing the recovery interval configuration option under the Server Properties - Database Settings page in SQL Server Management Studio. We can only change it to a value in minutes, and it will be used for all databases inside the SQL Server instance.

- **Manual.** You can manually cause checkpoints to occur by issuing the CHECKPOINT T-SQL command. Optionally, you can specify the time in seconds at which the checkpoint must be completed. If you do issue a manual checkpoint, it will run in the context of the current database. For example, executing CHECKPOINT 10 in a query window will perform a checkpoint within 10 seconds of the time you executed the query.
- **Indirect.** SQL Server 2012 added an extra option to configure checkpoint intervals on a per-database level. Configuring this option to a value greater than the default 0 will overwrite the automatic checkpoint process for the specific database. You can use indirect checkpoints for a specific database by using the following command: ALTER DATABASE [db name] SET TARGET_RECOVERY_TIME = [time in seconds or minutes].

As I mentioned before, SQL Server will attempt to throttle the checkpoint process to avoid overloading the storage subsystem if it believes this is necessary. It monitors the number of outstanding requests to the storage subsystem and tries to detect if there is any latency. Using this information, it will throttle the amount of IOs the checkpoint process generates so as to avoid a too-heavy load on the storage subsystem. When the checkpoint process is getting throttled, the SLEEP_BPOOL_FLUSH Wait Type will be recorded.

SLEEP_BPOOL_FLUSH Example

Generating SLEEP_BPOOL_FLUSH waits is relatively simple, and the script in Listing 5-10, which is almost the same one as we used for the LOGBUFFER and WRITELOG Wait Types, will put pressure on the checkpoint process such that SLEEP_BPOOL_FLUSH Waits will occur.

Listing 5-10. Generate SLEEP_BPOOL_FLUSH waits

```
USE trans_demo
GO

DECLARE @i INT
SET @i = 1

WHILE @i < 100

    BEGIN

        INSERT INTO transactions
            (t_guid)
        VALUES
            (newid())
```

```

SET @i = @i + 1

-- Force a checkpoint to occur within 1 second
CHECKPOINT 1

END

```

Since we are also using the same database as in the LOGBUFFER and WRITELOG Wait Types example, Listing 5-11 shows the script to create the database if it doesn't exist already.

Listing 5-11. Create trans_demo database

```

USE master
GO

-- Create demo database
CREATE DATABASE [trans_demo]
ON PRIMARY
(
    NAME = N'trans_demo', FILENAME = N'D:\Data\trans_demo.mdf' ,
    SIZE = 153600KB , FILEGROWTH = 10%
)
LOG ON
(
    NAME = N'trans_demo_log', FILENAME = N'D:\Log\trans_demo.ldf' ,
    SIZE = 51200KB , FILEGROWTH = 10%
)
GO

-- Make sure recovery model is set to full
ALTER DATABASE [trans_demo] SET RECOVERY FULL
GO

-- Create a simple test table
USE trans_demo
GO

CREATE TABLE transactions
(
    t_guid VARCHAR(50)
)
GO

```

What the script in Listing 5-10 will do is perform an insert of a random GUID into the transactions table inside a loop that is executed 100 times. Every time it enters a new GUID it will issue a CHECKPOINT command with a time limit of one second. This forces the checkpoint process to perform a checkpoint within the one-second time limit.

Before running the script in Listing 5-10, I cleared the `sys.dm_os_wait_stats` DMV using the `DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR)` command.

After almost 70 seconds the script completed on my test SQL Server. I then executed the following query to take a look at the `SLEEP_BPOOL_FLUSH` Wait Times:

```
SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'SLEEP_BPOOL_FLUSH';
```

The results of the query can be seen in Figure 5-29.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	SLEEP_BPOOL_FLUSH	311	60415	938	26

Figure 5-29. *SLEEP_BPOOL_FLUSH* waits

As you can see, the `SLEEP_BPOOL_FLUSH` Wait Time has a very high amount of Wait Time after running the script in Listing 5-10. Normally you would expect those Wait Times to be either very low or close to zero. If we were to remove the `CHECKPOINT` command from the script completely and let SQL Server decide on when to run the checkpoint process, we not only get a completely different result, as shown in Figure 5-30, but also the script's runtime is decreased to just a few milliseconds.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	SLEEP_BPOOL_FLUSH	0	0	0	0

Figure 5-30. *SLEEP_BPOOL_FLUSH* Wait Times after removing `CHECKPOINT`

Lowering `SLEEP_BPOOL_FLUSH` Waits

Even though it is not very common to run into performance problems caused by the `SLEEP_BPOOL_FLUSH` Wait Type, there are various methods to lower the Wait Times.

The most obvious one would be to check the various configuration options available to manually configure the recovery interval that we discussed earlier. The lower the value of the recovery interval, the more often checkpoint processes will take place, and the bigger the chance of running into `SLEEP_BPOOL_FLUSH` waits. Also, as you noticed in the example, performing frequent `CHECKPOINT` commands inside transactions can lead to `SLEEP_BPOOL_FLUSH` waits.

Another possible cause can be the storage subsystem on which your database data file resides. As explained earlier, the checkpoint process calculates the load of the storage subsystem and then decides if throttling its throughput is needed. If there is a frequent need of throttling because your storage subsystem is busy, you are more likely to see `SLEEP_BPOOL_FLUSH` waits occur.

SLEEP_BPOOL_FLUSH Summary

The SLEEP_BPOOL_FLUSH Wait Type is closely related to the checkpoint process in SQL Server. The checkpoint process is responsible for writing modified, or dirty, data pages from the buffer cache to the database data file. The checkpoint process analyzes the performance of the storage subsystem before it writes the dirty pages to disk, and if the storage subsystem is busy the checkpoint process will throttle its throughput, resulting in SLEEP_BPOOL_FLUSH waits. It is not very common to see very high SLEEP_BPOOL_FLUSH Wait Times, but they can impact performance nonetheless. Queries that frequently execute the CHECKPOINT T-SQL command, or a recovery interval that is configured to a very low value, can be possible causes for seeing SLEEP_BPOOL_FLUSH waits occur. The performance of your storage subsystem can also impact the checkpoint process if it is forced to throttle its throughput.

WRITE_COMPLETION

As with the ASYNC_IO_COMPLETION and IO_COMPLETION Wait Types, the WRITE_COMPLETION Wait Type is related to specific actions SQL Server performs on the storage subsystem. Again, it is very normal to see WRITE_COMPLETION waits occur on your SQL Server instance, and they should only be a cause for concern if the Wait Times are way higher than normal.

What Is the WRITE_COMPLETION Wait Type?

The WRITE_COMPLETION Wait Type is a relative of the IO_COMPLETION Wait Type. But where the IO_COMPLETION Wait Type is logged for specific read and write operations, the WRITE_COMPLETION Wait Type is only logged for some very specific write operations. Some of these write operations are growing a data or log file or performing the DBCC CHECKDB command.

Since the WRITE_COMPLETION Wait Type is related to writing SQL Server data to the storage subsystem, the performance of it can have an impact on the Wait Times.

WRITE_COMPLETION Example

To show you an example of a WRITE_COMPLETION wait occurring, I am going to perform a CHECKDB against the AdventureWorks database after clearing the sys.dm_os_wait_stats DMV using the DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR) command.

Keep in mind that this example is a completely normal situation in which WRITE_COMPLETION waits can occur, and it shouldn't stop you from performing regular database consistency checks!

Listing 5-12 shows the query I executed to generate a few WRITE_COMPLETION waits.

Listing 5-12. Generate WRITE_COMPLETION waits

```
DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR);

DBCC CHECKDB ('AdventureWorks');

SELECT * FROM sys.dm_os_wait_stats
WHERE wait_type = 'WRITE_COMPLETION';
```

The results of the last query in the batch are shown in Figure 5-31.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	WRITE_COMPLETION	2	1	1	1

Figure 5-31. WRITE_COMPLETION waits

As you can see, the amount of Wait Time is so low that it would not be a cause of any concern. This is also partly due to the fact that the AdventureWorks database is very small. Running CHECKDB against larger databases can result in higher Wait Times.

Lowering WRITE_COMPLETION Waits

If you see high WRITE_COMPLETION Wait Times, try to find out what process is generating the waits. In many cases it will be caused by a CHECKDB or database data or log file growth.

One thing worth checking is the instant file initialization option discussed in the ASYNC_IO_COMPLETION section earlier in this chapter. Not using this option can impact the duration of the WRITE_COMPLETION Wait Time.

Another, far less common cause for a higher WRITE_COMPLETION Wait Time is when you are experiencing page latch contention on your Page Free Space page (or PFS). The PFS page tracks the amount of free space in data pages. If a process needs to modify the PFS page very frequently, it is possible to see WRITE_COMPLETION waits occur along with many PAGELATCH_UP waits, which we will discuss in Chapter 8. To give you an example of such a scenario, consider a high amount of concurrent queries that all create a temporary table, insert a few rows, and remove the temporary table again. In this case the PFS page of the tempdb database needs to get updates very frequently to reflect the creation and removal of the temporary tables.

WRITE_COMPLETION Summary

The WRITE_COMPLETION Wait Type, just like the ASYNC_IO_COMPLETION and IO_COMPLETION Wait Types, is related to specific storage-related actions performed by SQL Server. Seeing WRITE_COMPLETION waits is very normal and won't be cause for concern in many situations. Operations such as CHECKDB and database data or log file growth can cause WRITE_COMPLETION waits.



Backup-Related Wait Types

Backups are a very important part of database administration, and in many cases they are essential for the survival of the company you work for. Data has become so important for businesses that if some disaster causes data to be lost, companies can lose large amounts of money, or even go out of business.

There are many methods we can implement on all levels of your IT infrastructure to make sure no data is lost (or as little as possible) during a disaster. We could implement a SAN to make sure our data is not stored on a local disk drive on your server. Or we could design SQL Server AlwaysOn Availability Groups to replicate our data across datacenters. But the first step we should take, and hopefully have already taken, is performing regular backups of our data inside our SQL Server databases.

Implementing and scheduling SQL Server backups is not a very difficult task, and there is no excuse not to perform a backup. The type of backup and the interval of backup operations are dictated by the needs of the organization you work for and are frequently expressed in “RTO” (Recovery Time Objective) and “RPO” (Recovery Point Objective) times. These times represent the amount of time it should take to recover from a disaster, and the amount of data loss that is acceptable when a disaster occurs. These two times should be the primary input for your SQL Server backup strategy.

Thankfully, SQL Server has different options available to us for meeting RTO and RPO requirements right out of the box. This means we can use SQL Server’s own backup mechanism to fulfill our company’s RTO and RPO times; we are not necessarily dependent on third-party backup software. Since the SQL Server backup operation is an internal process, there are different Wait Types associated with it, and in this chapter we will take a look at three of the most common Wait Types that are directly related to performing backups and restores.

Noticing high Wait Times on these backup/restore-related Wait Types will not likely lead to a performance degradation of your SQL Server instance. However, we do have options to optimize the SQL Server backup process that can result in faster backup and restore times. And since backups/restores of your database(s) are vital for the survival of your company, optimizing backup and restore throughput can be well worth the effort.

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-1-4842-1139-7_6](https://doi.org/10.1007/978-1-4842-1139-7_6)) contains supplementary material, which is available to authorized users.

BACKUPBUFFER

The first backup-related Wait Types we will discuss is BACKUPBUFFER. If we look up the definition of this Wait Type on Books Online we would get the following text: “Occurs when a backup task is waiting for data, or is waiting for a buffer in which to store data. This type is not typical, except when a task is waiting for a tape mount.” Apparently, we would only see this Wait Type when we are writing our backups to a tape device, and this is wrong. BACKUPBUFFER waits will practically always be logged during a backup operation, no matter the destination of the backup file. The reason for this is in the way the SQL Server backup operation uses buffers to read data from the database and write it to the backup file.

What Is the BACKUPBUFFER Wait Type?

To understand how BACKUPBUFFER waits are generated we have to take a look at the internals of the SQL Server backup process. These internals are mostly the same regardless of the backup method you use (i.e., transaction log, differential or full backup) and as such they will encounter the same Wait Types.

SQL Server allocates buffers for the backup process. These buffers will be filled with data from your database and will be moved through a backup/restore process in order to get written to the backup file (or vice versa for a restore operation). The buffers are allocated inside the memory of your system, but outside the memory of your buffer cache so as to avoid stealing memory from the buffer cache. The size and the amount of the backup buffers are automatically calculated by SQL Server, but we can configure these values ourselves as parameters of the backup/restore command. Figure 6-1 shows how these backup buffers are ordered and moved through a “reader,” which reads the data from your database or backup file to a buffer, and a “writer,” which writes the data from the buffer to the backup file or database.

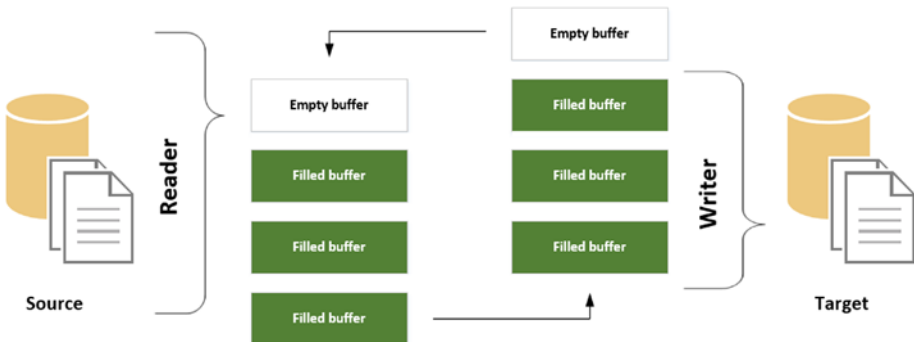


Figure 6-1. Backup buffers moving through reader and writer

We can view information about the buffer amount and size during a backup or restore operation by enabling two trace flags, 3213 and 3605, which will output backup/restore information into the SQL Server error log. The query in Listing 6-1 enables both trace flags and performs a full database backup of the AdventureWorks database on my test SQL Server.

Listing 6-1. Full database backup with backup-information trace flags

```
-- enable trace flags
DBCC TRACEON (3213);
DBCC TRACEON (3605);

-- backup database
BACKUP DATABASE [AdventureWorks]
  TO DISK = N'F:\Backup\aw_21042015.bak'
  WITH NAME = N'AdventureWorks-Full Database Backup';
GO

-- disable trace flags
DBCC TRACEOFF (3213);
DBCC TRACEOFF (3605);
```

Keep in mind that trace flags inside SQL Server should only be used under the guidance of Microsoft Support. I am enabling them now to show me backup information on my test SQL Server, but I would advise against using them on a production system.

Inside the SQL Server error log, additional information about the backup we just performed is logged, as you can see in Figure 6-2.

4/21/2015 10:56:57 AM	spid59	DBCC TRACEOFF 3605, server process ID (SPID) 59. This is an informational message only; no user action is required.
4/21/2015 10:56:57 AM	spid59	DBCC TRACEOFF 3213, server process ID (SPID) 59. This is an informational message only; no user action is required.
4/21/2015 10:56:57 AM	Backup	Database backed up. Database: AdventureWorks2012, creation datetime): 2015/04/03/10:51:17, pages dumped: 162
4/21/2015 10:56:47 AM	spid59	Media Buffer size: 1024KB
4/21/2015 10:56:47 AM	spid59	Media Buffer count: 7
4/21/2015 10:56:47 AM	spid59	Filesystem i/o alignment: 512
4/21/2015 10:56:47 AM	spid59	TXF device count: 0
4/21/2015 10:56:47 AM	spid59	Filestream device count: 0
4/21/2015 10:56:47 AM	spid59	Fulltext data device count: 0
4/21/2015 10:56:47 AM	spid59	Tabular data device count: 1
4/21/2015 10:56:47 AM	spid59	Total buffer space: 7 MB
4/21/2015 10:56:47 AM	spid59	Min MaxTransferSize: 64 KB
4/21/2015 10:56:47 AM	spid59	MaxTransferSize: 1024 KB
4/21/2015 10:56:47 AM	spid59	Sets Of Buffers: 1
4/21/2015 10:56:47 AM	spid59	BufferCount: 7
4/21/2015 10:56:47 AM	spid59	Memory limit: 127MB
4/21/2015 10:56:47 AM	spid59	Backup/Restore buffer configuration parameters
4/21/2015 10:56:47 AM	spid59	DBCC TRACEON 3605, server process ID (SPID) 59. This is an informational message only; no user action is required.
4/21/2015 10:56:47 AM	spid59	DBCC TRACEON 3213, server process ID (SPID) 59. This is an informational message only; no user action is required.

Figure 6-2. Additional backup information

In this case, the backup operation created seven buffers, shown by the `BufferCount` parameter, with a size of 1024 KB each, as shown by the `MaxTransferSize` parameter. The total memory needed to create the buffers is shown by the `Total_buffer_space` parameter, 7 MB (`BufferCount * MaxTransferSize`). Another interesting bit of information that is returned is the `memory_limit`. This will show the maximum amount of memory outside of the buffer cache that the backup operation could access.

Now that we have an idea of how the backup process works inside SQL Server, let's take a look where the `BACKUPBUFFER` Wait Type comes in.

As we described earlier, the SQL Server backup process uses buffers to store data that needs to be written to the backup file. Whenever a buffer is not directly available, the `BACKUPBUFFER` wait will occur, making the process wait until a full buffer is written to the backup file and it becomes available again.

BACKUPBUFFER Example

Generating `BACKUPBUFFER` waits is very simple—just perform a backup operation. For this example I ran the query shown in Listing 6-2. The query will first reset the `sys.dm_os_wait_stats` DMV, then will perform a full backup of the AdventureWorks database, and finally will return the Wait Statistics information for the `BACKUPBUFFER` Wait Type.

Listing 6-2. Generating `BACKUPBUFFER` waits

```
-- clear sys.dm_os_wait_stats
DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR);

-- backup database
BACKUP DATABASE [AdventureWorks]
  TO DISK = N'F:\Backup\aw_21042015.bak'
WITH
  NAME = N'AdventureWorks-Full Database Backup';
GO

-- Query BACKUPBUFFER waits
SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'BACKUPBUFFER';
```

The results of the query against the `sys.dm_os_wait_stats` DMV are shown in Figure 6-3.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	BACKUPBUFFER	177	292	11	24

Figure 6-3. `BACKUPBUFFER` waits

The total duration of the backup operation was around 10 seconds on my test SQL Server. Of those 10 seconds, only 292 milliseconds were spent waiting on free backup buffers.

Lowering BACKUPBUFFER Waits

As stated in the introduction of this chapter, backup-related waits aren't normally any cause for concern since they normally won't impact the performance of your SQL Server instance. However, we can improve backup performance by using the Wait Statistics information of the various backup-related Wait Types.

One of the most common ways to lower BACKUPBUFFER Wait Times is by adding more buffers for the backup operation to use, overwriting the automatic allocation of buffers. We can do this by specifying the BUFFERCOUNT option inside the BACKUP T-SQL command. There is, however, a catch to altering the number of buffers the backup operation can use. Every buffer created will allocate the value of the MAXTRANSFERSIZE option; this value can be automatically calculated by SQL Server itself or by setting the value yourself inside the BACKUP command (up to a maximum of 4194304 bytes). Since the backup operation allocates memory outside of the buffer cache, there is a chance that using too many or too large buffers can result in out-of-memory problems. So, be careful when testing what the optimal value for your SQL Server instance is.

Listing 6-3 shows a modification of the query in Listing 6-2, which we used to demonstrate BACKUPBUFFER waits occurring. In this case we added the BUFFERCOUNT option and configured it to a value of 200.

Listing 6-3. Database backup with BUFFERCOUNT configured

```
-- clear sys.dm_os_wait_stats
DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR);

-- backup database
BACKUP DATABASE [AdventureWorks]
  TO DISK = N'F:\Backup\aw_21042015.bak'
WITH
  NAME = N'AdventureWorks-Full Database Backup',
  BUFFERCOUNT = 200;
GO

-- Query BACKUPBUFFER waits
SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'BACKUPBUFFER';
```

The results of the query against the `sys.dm_os_wait_stats` DMV are shown in Figure 6-4.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	BACKUPBUFFER	194	115	8	6

Figure 6-4. BACKUPBUFFER waits

As you can see, the amount of time spent on the BACKUPBUFFER Wait Time went down to 115 milliseconds instead of the 292 milliseconds it spent when we did not supply the BUFFERCOUNT parameter.

Another option is to configure the MAXTRANSFERSIZE option inside the BACKUP T-SQL command. This will allow buffers to be filled with larger units of works, up to a value of 4194304 bytes, or 4 MB. Again, allocating more space for the buffers will result in a larger reservation of memory.

BACKUPBUFFER Summary

BACKUPBUFFER waits occur normally during backup or restore operations when the backup/restore operation has to wait for free buffers to become available again. Because they occur normally they shouldn't be a cause for concern. We do have some options for lowering BACKUPBUFFER Wait Times that will also impact the duration of the backup/restore operation. They should be configured and tested thoroughly though, because setting those parameters too high can result in out-of-memory errors.

BACKUPIO

Just like the BACKUPBUFFER Wait Type, the BACKUPIO Wait Type occurs when a part of the backup or restore operation runs into contention problems. Another similarity is the description of this Wait Type on Books Online: "Occurs when a backup task is waiting for data, or is waiting for a buffer in which to store data. This type is not typical, except when a task is waiting for a tape mount." Again, this Wait Type is common when performing a backup or restore operation, even when the backup target, or restore source, is not a tape device.

What Is the BACKUPIO Wait Type?

To better understand how BACKUPIO waits are generated, we have to take a look at Figure 6-5, which we showed earlier as Figure 6-1.

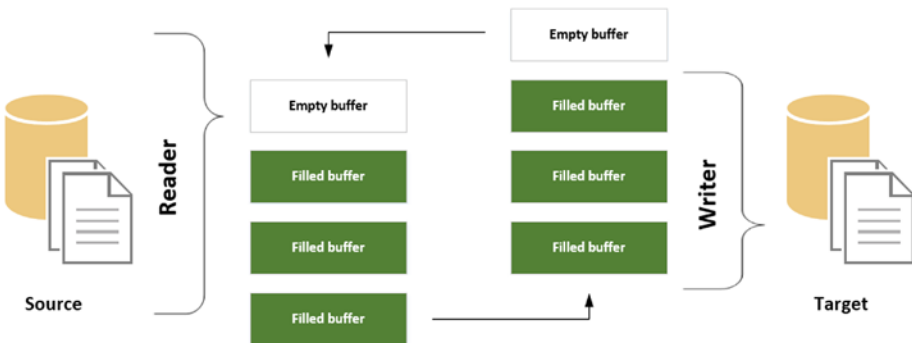


Figure 6-5. Internals of a backup operation

In the previous section where we discussed the BACKUPBUFFER Wait Type we explained that the BACKUPBUFFER Wait Type occurs when we are waiting for a free (empty) buffer to become available. For the most part, the BACKUPBUFFER Wait Type is situated on the left side of Figure 6-5, at the reader. The BACKUIO Wait Type occurs for the most part on the right side of Figure 6-5, at the writer section. When BACKUIO waits occur there is a delay in the time the writer is writing data. This delay can be caused by many different things; for instance, when writing a backup to a slow disk, writing a backup to a network location, or when restoring a database.

The BACKUIO Wait Type will frequently be accompanied by ASYNC_IO_COMPLETION waits when a database backup or restore is performed.

BACKUIO Example

We can make use of the same example as we used to demonstrate the BACKUPBUFFER Wait Type. I did modify the query a little to return BACKUIO waits instead of BACKUPBUFFER waits, and I also included the ASYNC_IO_COMPLETION in the results of the query against the sys.dm_os_wait_stats DMV. Listing 6-4 shows the modified backup query.

Listing 6-4. Generating BACKUIO waits

```
-- clear sys.dm_os_wait_stats
DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR);

BACKUP DATABASE [AdventureWorks]
  TO DISK = N'F:\Backup\aw_21042015.bak'
WITH
  NAME = N'AdventureWorks-Full Database Backup';
GO

-- Query BACKUIO waits
SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'BACKUIO'
OR wait_type = 'ASYNC_IO_COMPLETION';
```

The results of the query against the sys.dm_os_wait_stats DMV can be seen in Figure 6-6.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	ASYNC_IO_COMPLETION	3	9577	9402	0
2	BACKUIO	1555	4259	31	56

Figure 6-6. ASYNC_IO_COMPLETION and BACKUIO waits

As you can see in Figure 6-6, the database backup caused both Wait Types to be generated, and the most time has been spent on the `ASYNC_IO_COMPLETION` wait, which is responsible for reading the data pages that need to be written to the backup file. Since my backup destination is on an SSD disk we didn't encounter very high `BACKUIO` Wait Times.

Lowering BACKUIO Waits

Tweaking the `BUFFERCOUNT` and `MAXTRANSFERSIZE` options do not have as much impact on the `BACKUIO` Wait Type as they did on the `BACKUPBUFFER` Wait Type. When you see higher than normal Wait Times on the `BACKUIO` Wait Type, the problem is most likely related to the throughput of either your storage subsystem or network location you are writing or reading your backup to/from. Make sure to check both locations for possible performance problems like high latency or network utilization.

BACKUIO Summary

Just like the `BACKUPBUFFER` Wait Type, the `BACKUIO` Wait Type occurs when a backup or restore operation is being performed. While the `BACKUPBUFFER` Wait Type is mostly related to the speed at which the backup operation can access the backup buffers, the `BACKUIO` Wait Type is related to the speed at which those backup buffers can be written to disk. `BACKUIO` waits frequently occur together with `ASYNC_IO_COMPLETION` waits when performing full database backups or restores. When seeing higher than normal Wait Times for the `BACKUIO` Wait Type, check the performance metric of the location you are writing or reading the backup file to or from. Lowering `BACKUIO` Wait Times will not have an impact on the query performance of your system, but will help speed up backup and restore operations.

BACKUPTHREAD

The `BACKUPTHREAD` Wait Type is frequently seen when performing restore operations on a database, but can also occur during a backup operation. It occurs when another thread is waiting for the backup/restore operation to finish so it can continue processing.

What Is the BACKUPTHREAD Wait Type?

When you see `BACKUPTHREAD` waits occurring it means that another thread wants to access a resource that is currently being accessed by a backup or restore operation. During the time the thread has to wait for the backup/restore to complete, `BACKUPTHREAD` Wait Time will be recorded. An example of this type of wait would be a thread that wants to access the database data file while it is being restored; for instance, the `ASYNC_IO_COMPLETION` Wait Type that is writing the data file to disk.

`BACKUPTHREAD` waits are not usually a cause for concern. They only indicate that other threads are waiting for the backup/restore operation to complete, and they frequently have the same duration as the time it took for your backup or restore to complete. They do, however, give you a hint that there are other waits occurring that might deserve investigation if the Wait Times are higher than expected.

Because a picture says more than a thousand words, Figure 6-7 shows the relation of the BACKUPTHREAD Wait Type with a restore operation and other waits occurring.




Figure 6-7. BACKUPTHREAD relation to other threads

In Figure 6-7 you can see that the BACKUPTHREAD wait is occurring because another thread also wanted to access a resource that was currently owned by the restore operation.

BACKUPTHREAD Example

An easy way to demonstrate BACKUPTHREAD waits occurring is by performing a restore operation. When you perform a restore other processes will need to access the database data files to write the information from the backup file to the database data files.

Listing 6-5 shows a script to restore a backup file I made earlier of the AdventureWorks database on my test SQL Server.

Listing 6-5. Restore AdventureWorks database

```
-- Restore database
USE [master]
RESTORE DATABASE [AdventureWorks]
FROM DISK = N'F:\Backup\AWBackup.bak'
WITH FILE = 1, REPLACE;
GO
```

If we were to look at the `sys.dm_os_waiting_tasks` DMV while the backup is running, we would see the waits occurring as shown in Figure 6-8, which was taken on my test SQL Server.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address
1	0x000000017F266928	64	1	8675	PREEMPTIVE_OS_WRITEFILEGATHER	NULL
2	0x000000017F02DC38	64	0	8649	BACKUPTHREAD	NULL
3	0x00000001672AC928	64	3	9	BACKUPIO	NULL

Figure 6-8. BACKUPTHREAD and other waits

As you can see in Figure 6-8, the Wait Time of the BACKUPTHREAD Wait Type is almost identical to that of the PREEMPTIVE_OS_WRITEFILEGATHER Wait Type. This Wait Type is responsible for writing data to the file system, but we will dive deeper into this specific Wait Type in Chapter 10, “Preemptive Wait Types.”

Lowering BACKUPTHREAD Waits

While the BACKUPTHREAD Wait Type itself doesn’t indicate any problems, its combination with other Wait Types can be a reason for some additional research. Basically, every method you can use to speed up your backup or recovery process will have an impact on the BACKUPTHREAD Wait Time.

Some good pointers to start with are the BufferCount and MaxTransferSize options that you can specify on the BACKUP and RESTORE T-SQL commands. We touched upon these settings when we discussed the BACKUPBUFFER and BACKUPIO Wait Types. Tweaking these settings can make your backups and restores take less time, resulting in lower BACKUPTHREAD Wait Times.

Another setting that can dramatically improve backup and restore times is the instant file initialization option that we discussed in Chapter 5, “IO Related Wait Types,” in the ASYNC_IO_COMPLETION section.

BACKUPTHREAD Summary

The BACKUPTHREAD Wait Time doesn’t indicate access to a specific resource, but rather indicates that another process is waiting for a backup or restore operation to complete. It is very common to see this Wait Type, especially during restore operations. Lowering the duration of backup and restore operations will also be reflected in the Wait Times of the BACKUPTHREAD Wait Type. One of the methods you can use to lower BACKUPTHREAD Wait Times is checking whether instant file initialization is enabled. This setting does not directly impact the BACKUPTHREAD Wait Type, but it will impact other Wait Types, which will in turn impact the BACKUPTHREAD Wait Type.

CHAPTER 7



Lock-Related Wait Types

Locking is a fundamental part of every relational database, or Relational Database Management System (RDBMS). SQL Server is based on the relational database model, and as such uses locking when data is accessed. Even though we frequently relate locking to performance problems, it plays a vital role in making sure your data is reliable during concurrent workloads. The way SQL Server, or any other RDBMS for that matter, takes care of this data reliability is by following the “ACID” properties, which were originally defined by Jim Gray in the 1970s but received their name in 1983 from Andreas Reuter and Theo Härder. These ACID properties are enforced upon single operations, which we know as transactions. The acronym ACID consists of four characteristics that guarantee data reliability inside transactions. The list below describes each of these characteristics.

- **Atomicity.** The Atomicity characteristic requires that transactions are all or nothing. This means that if one part of the transaction fails, the complete transaction fails, and every change done inside the transaction needs to be changed back to the state before the transaction started.
- **Consistency.** The Consistency characteristic requires that data written to the database by the transaction is legal. This means that the data must be stripped of illegal or bad input.
- **Isolation.** The Isolation characteristic requires that every transaction is hidden from other concurrent transactions. From a transaction point of view, this means every transaction is executed serially.
- **Durable.** The Durable characteristic requires that every committed transaction remains committed, even in the event of a power failure or disaster.

As you might have guessed from reading the different ACID properties, locking inside SQL Server is closely related to the Isolation characteristic.

Since this chapter is dedicated to lock-related Wait Types, we won't go into detail about ACID properties besides Isolation. If you are interested in learning more about the

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-1-4842-1139-7_7](https://doi.org/10.1007/978-1-4842-1139-7_7)) contains supplementary material, which is available to authorized users.

ACID properties and database theory, a good place to start would be the “Principles of Transaction-Oriented Database Recovery” research paper by Andreas Reuter and Theo Harder, which describes the ACID properties in detail.

To get a better understanding of how the Isolation characteristic works, we need to understand transactions. A transaction represents an interaction with the database that can consist of multiple actions and that is separated from other transactions.

To make sure our transactions do not conflict with other, concurrent, transactions, SQL Server uses locks. These locks make sure no other transaction can modify data that your transaction is processing at the same time. For example, if you make a \$100 withdrawal from your bank account, you do not want another, concurrent, withdrawal to modify that amount. Other transactions will have to wait with their withdrawals until your transaction is completed. Inside SQL Server, this process works the same way. When you request data from your database you want to get the data returned you asked for, without running the risk that the data is being modified while you are requesting it.

When you run your transaction it will be protected by a lock SQL Server places on the object you are accessing. If another transaction wants to interact with the same object, a block will occur. When this block occurs, the latter transaction will have to wait until the lock on the object is removed. The transaction can then place its own lock on the object and start its interaction.

There are many options available to us within SQL Server to control the behavior of locking and blocking, and most of them are related to changing the Isolation of certain or all transactions against a database. There is also a lot of information about locking and blocking we can access inside SQL Server, and not the least of these are inside Wait Statistics. The time a transaction is waiting to access a locked object is recorded as Wait Time for specific, lock-related Wait Types (depending on the type of lock the transaction intends to place).

In this chapter we will discuss the various Wait Types that are related to locking and blocking and how we can lower or even resolve them. This requires some knowledge of how SQL Server uses locks, and for this reason I included a section to familiarize ourselves with locking and blocking before we dive into the lock Wait Types.

Introduction to Locking & Blocking

As we just discussed, SQL Server uses locks to isolate different concurrent transactions from each other so data is only accessed or modified by one transaction at a time. There are different lock types, or lock modes, SQL Server can use, and there are various object levels SQL Server can place locks on. To make it even more complex, different lock modes are not necessarily compatible with each other, and when two incompatible locks meet, a block occurs.

Lock Modes and Compatibility

To start off, let's get ourselves familiar with the different types of locks, or lock modes, inside SQL Server. The list that follows describes the most common lock modes. There are more lock modes inside SQL Server, but those only occur when you perform very specific actions. A complete list of the different lock modes can be found on the MSDN page that discusses lock modes here: <https://technet.microsoft.com/en-us/library/ms175519.aspx>. SQL Server uses acronyms to indicate which lock mode is being used inside SQL Server. These acronyms are shown in parentheses.

- **Shared (S).** A Shared lock will be placed on a resource when a query is selecting data from that resource. For instance, a `SELECT * FROM [table]`.
- **Update (U).** The Update lock mode is used when a query wants to modify a resource. It was introduced to prevent “deadlocks,” a situation where locks are waiting on each other to release in concurrent transactions that want to modify the same resource.
- **Exclusive (X).** An Exclusive lock is placed when a transaction wants to modify the resource. When an Exclusive lock is in place, no other transactions can modify the resource. For instance, `INSERT`, `UPDATE`, or `DELETE` T-SQL statements will result in Exclusive locks.
- **Schema (Sch).** Schema locks are used when a table is being modified. An example of this would be adding a column to a table.
- **Intent (I).** Intent locks are used to indicate that locks are placed at a lower level in the locking hierarchy. We will go into more detail on the lock hierarchy in a bit.

When different locks need to interact with each other, SQL Server performs a lock compatibility check on the different lock modes involved. Not all of the lock modes are compatible with each other, which means that when two different transactions are not able to access the resource at the same time because of incompatible locks, a block will occur. For instance, when a Shared lock is placed to read from a row, and another transaction wants to modify the row by placing an Exclusive lock, the Exclusive lock will have to wait until the Shared lock is removed. Table 7-1 shows the lock mode compatibility for the Shared, Update, and Exclusive lock modes.

Table 7-1. Lock Compatibility

Lock mode	Shared	Update	Exclusive
Shared	Yes	Yes	No
Update	Yes	No	No
Exclusive	No	No	No

Let’s go through an example to illustrate lock compatibility. Say you want to read from a row inside a table by executing a `SELECT` statement against that table. When you execute your query SQL Server will check if there is any existing lock already in place on the row you want to access, and if it is compatible with the lock you want to place on the row. Let’s assume there isn’t a lock in place when you run your query. In this case a Shared lock will be placed on the row, indicating that your query is reading data from that row. Right after you execute your query another transaction is issued by another user that wants to modify data inside the row you are accessing. SQL Server will detect that there already is a Shared lock in place on the row, making the second transaction wait before placing its Exclusive

lock, since Shared and Exclusive locks are incompatible. The user who ran the second transaction might experience a delay, since the transaction is waiting for the Shared lock to be removed before its Exclusive lock can be placed. If a third transaction is started that wants to read the same row as your transaction, no lock conflict will occur. Shared locks are compatible with other Shared locks, meaning that the third transaction does not have to wait to place its lock, and it directly receives the results it asked for.

Figure 7-1 shows the example above, where the dotted line indicates an incompatible lock that has to wait.

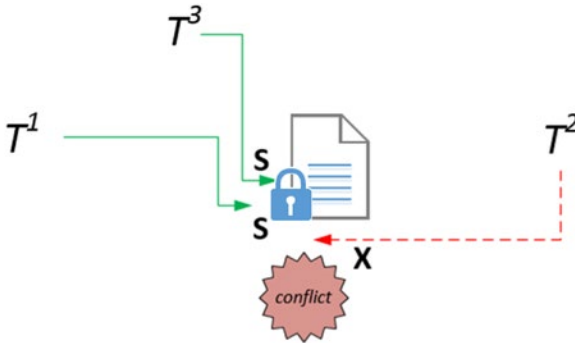


Figure 7-1. Concurrent lock situation

Locking Hierarchy

SQL Server uses multigranular locking to allow different locks for different-level objects. It does this to minimize the overhead cost of locking. The lowest possible object where a lock can be placed is a row, and the largest is the database. There are many levels between those two granularity levels, and SQL Server automatically decides on what level the lock should be placed to minimize locking overhead. The following list shows the most common lock levels, ordered from the highest granularity to the smallest:

- Database
- Database file
- Table/Object
- Extent
- Page
- RID (row inside a heap) / KEY (row inside a clustered index)

The Intent locks we discussed earlier also play an important part in the placement of locks upon the different granularity levels. SQL Server will place Intent locks on objects that are on a higher granularity to indicate a lock has been placed at a lower level. This protects the lower-level locks from changes on objects at a higher granularity level. All the Intent locks that are placed, from the highest granularity level to the actual lock on an object, when looked at together are called the locking hierarchy.

Figure 7-2 shows a graphical representation of a locking hierarchy for the modification of data inside a row, which will require an Exclusive lock on the row and Intent Exclusive locks higher in the hierarchy.

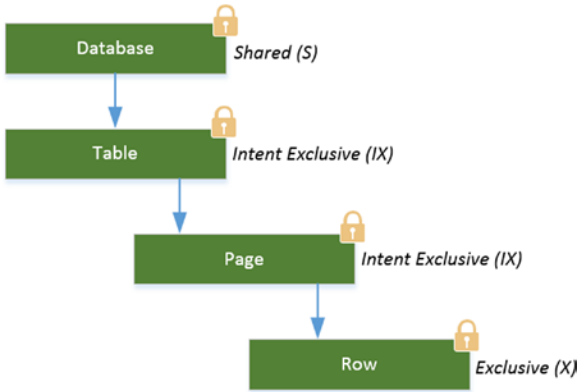


Figure 7-2. Lock hierarchy example

Note the Shared lock on the database level. Every request will always place one to protect changes to the database while transactions are active. This makes sure that, for instance, you cannot delete a database while transactions are still active. Also note that the Intent locks will use the same lock mode on the lowest object, in this case Intent Exclusive (IX). If a Shared lock was placed, the lock mode of the Intent lock would change as well, in this case to Intent Shared (IS). We will go a little deeper into Intent locks a bit further on in this chapter.

Isolation Levels

We can exercise a certain level of control over what locks are being placed by a transaction by changing the Isolation level. The Isolation level defines the degree to which transactions are isolated from each other during concurrent operations. We can change the Isolation level on either a connection or a transaction basis. Changing the Isolation level will only change the behavior of Shared locks; Exclusive locks that are needed for data modification are not affected. Changing the Isolation level will also introduce certain phenomena. These phenomena have an impact on the results of your read transaction and occur because of the changes to how Shared locks are placed and held during the transaction. The list that follows shows the various Isolation levels, from the lowest form of Isolation to the highest, available in SQL Server and the phenomena related to them:

- **Read Uncommitted.** This Isolation level will allow reads to occur while another transaction is performing modifications on the same object. It will not wait until the Exclusive lock on the object is released. This makes it possible to read uncommitted values

called “dirty reads.” Dirty reads can be bad (if you do not expect them) because they can return a value that is no longer current in the database. For instance, if someone is updating a value to “B” while it was “A” at the start of the transaction, other users that query the same data at the same time can get the old value of “A” back instead of the updated “B” value.

- **Read Committed.** This is the default Isolation level in SQL Server. Using this Isolation level will make read transactions wait until concurrent write transactions are completed. A Shared lock will be placed on a row and will be released right after the row has been read. The phenomena associated with this Isolation level is called “inconsistent analysis.” This means that it is possible to receive different results from the same read query if the data were modified by another transaction in the time between both read transactions.
- **Repeatable Read.** Setting the Isolation level to Repeatable Read will lock rows that are being read by a transaction. But instead of releasing the Shared lock on the row after it has been read, Repeatable Read will keep the lock in place until the entire transaction is completed. A Repeatable Read makes it possible for “phantom reads” to occur. Phantom reads occur whenever data is added or changed by another transaction that has not yet been locked by the read transaction.
- **Serializable.** The Serializable Isolation level is the highest possible Isolation level you can use, and that means it will place the most locks to ensure the data you are reading is not modified during the time the transaction is running. It does this by locking the entire range of data (for instance, an entire table) you are selecting, making it impossible to make changes to that data. Since the entire range of data you are selecting is being locked right at the start of the transaction, there are no phenomena possible.

SQL Server 2005 added another method for isolating transactions called Row Versioning. Row Versioning uses versions of data modification and returns them to read queries without causing blocking. When a transaction modifies data, that change will be recorded as a version. When a read transaction accesses the same data it will receive the version of the change before the modification transaction is committed. More information about Row Versioning can be found on Books Online at <https://technet.microsoft.com/en-us/library/ms189050.aspx>.

Because Isolation levels, and their locking behavior, can be complex to understand, I added Figure 7-3, which shows the way the various Isolation levels implement locking during a read operation. The boxes represent rows inside a table, and a row with a lock means a Shared lock is active on that row.

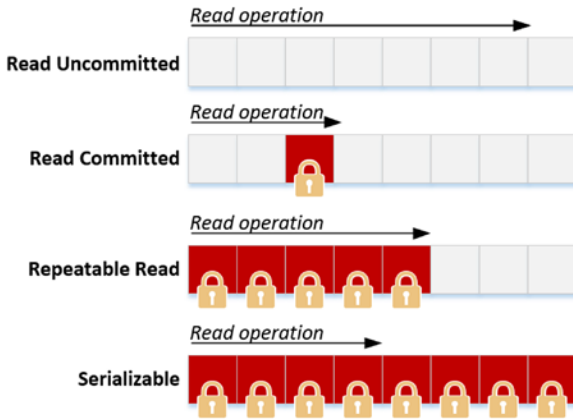


Figure 7-3. Isolation levels and locking behavior

There are various reasons why you would want to use a different Isolation level than the default of Read Committed. In many cases these reasons are related to the amount of locking/blocking you expect with your workload, or how “correct” the data returned by your transaction should be. For instance, with the default Isolation level of Read Committed it is possible that data is modified by other transactions while your transaction is running, which means that the results at the end of the transaction are not the same as they were at the start of your transaction. To make sure no data can change while your transaction is running, you could use the Serializable Isolation level, but this means more locks need to be placed and maintained, resulting in more blocking in concurrent SQL Server environments.

We can only change the default Isolation level of Read Committed by specifically configuring a different Isolation level for a connection or by supplying a table hint (an exception is Snapshot Isolation, which is configured at the database level). For instance, the two queries that follow show two different methods of executing a query using the Read Uncommitted Isolation level. The first query sets the transaction Isolation level for the entire session:

```
SET TRANSACTION ISOLATION LEVEL READ UNCOMMITTED
GO

BEGIN TRANSACTION

SELECT *
FROM Person.Person

COMMIT TRANSACTION;
GO
```

Another method is to use a table hint to set the Isolation level to Read Uncommitted:

```
SELECT *
FROM Person.Person
WITH (READUNCOMMITTED);
```

Both of these methods will achieve the same effect, but keep in mind that setting the Isolation level for the session will result in using the selected Isolation level for all the queries that are being executed in this specific session after setting it.

Querying Lock Information

To take a look at currently placed locks we can use the `sys.dm_tran_locks` DMV. This DMV will return a row for every active lock inside the SQL Server instance, along with information like the type of lock, the resource type, the session ID that placed the lock, and whether the lock is granted or is waiting to be placed. Figure 7-4 shows a (small) portion of the output of the DMV on my test SQL Server machine.

	resource_type	resource_subtype	resource_database_id	resource_description	resource_associated_entity_id	resource_lock_partition	request_mode	request_type	request_status
51	KEY		5	(93c9639b6eed)	7205759404555648	0	X	LOCK	GRANT
52	KEY		5	(2c515590e3e)	72057594227982336	0	X	LOCK	GRANT
53	PAGE		5	1:10672	72057594046971904	0	IX	LOCK	GRANT
54	PAGE		5	1:10672	72057594046971904	0	S	LOCK	WAIT
55	KEY		5	(fc98b13f12e)	72057594227982336	0	X	LOCK	GRANT
56	KEY		5	(6174a3a2309)	72057594046971904	0	X	LOCK	GRANT
57	KEY		5	(60a92478296)	72057594227982336	0	X	LOCK	GRANT

Figure 7-4. `sys.dm_tran_locks` output

If we take a look at Figure 7-4 we can see that a number of Exclusive locks (X) have been granted and are placed at the Key lock level. This means a transaction is currently modifying data inside a clustered index. There is also an Intent Exclusive lock on the Page level, which is above the Key lock level, indicating that there is an Exclusive lock lower down in the hierarchy. Also note that a Shared lock is currently waiting to get placed on the same data page (1:10672). The lock cannot be granted just yet, as there is an incompatible Intent Exclusive lock in place.

Since the Shared lock has to wait before it can be placed on the data page, we can view the time it has been waiting by looking at the Wait Statistics. Figure 7-5 shows a part of the results of a query against the `sys.dm_os_waiting_tasks` DMV.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address	blocking_task_address	blocking_session_id
5	0x000000017f01d0c3	15	0	149555910	CHECKPOINT_QUEUE	0x0000000037f6ec00	NULL	NULL
6	0x000000017f01d498	5	0	3567	REQUEST_FOR_DEADLOCK_SEARCH	0x000000003827eecd	NULL	NULL
7	0x000000017f26e558	9	0	176332298	BROKER_TRANSMITTER	NULL	NULL	NULL
8	0x000000017f267c38	58	0	402907	LCK_M_S	0x000000016fce3280	NULL	57
9	0x000000017f024cf8	2	0	117	LAZYWRITER_SLEEP	NULL	NULL	NULL
10	0x000000017f024928	12	0	231616	SP_SERVER_DIAGNOSTICS_SLEEP	0x0000000000000001	NULL	NULL

Figure 7-5. Lock information inside `sys.dm_os_waiting_tasks`

By using the `sys.dm_os_waiting_tasks` DMV, we can see that session ID 58 is currently waiting on a resource named `LCK_M_S`. This represents a Shared lock resource type. Session ID 58 is currently being blocked by session ID 57, which happens to be the same session that has the Exclusive and Intent Exclusive locks placed on the objects session ID 58 is trying to query. The `sys.dm_os_waiting_tasks` DMV will also return information we can use as input for the `sys.dm_tran_locks` DMV. This information will be available in the `resource_description` column of the `sys.dm_os_waiting_tasks` DMV, as shown in Figure 7-6.

blocking_exec_context_id	resource_description
NULL	NULL
NULL	NULL
NULL	NULL
NULL	pagelock fileid=1 pageid=10672 dbid=5 subresource=FULL id=lock16fd27700 mode=IX associatedObjectId=72057594046971904
NULL	NULL
NULL	NULL

Figure 7-6. *resource_description* column of the `sys.dm_os_waiting_tasks` DMV during a block

If we copy the `associatedObjectId` and use it as input in the `WHERE` clause against the `sys.dm_tran_locks` DMV, we will receive more information about why, and on what, this task is waiting. The following query will retrieve all the rows inside the `sys.dm_tran_locks` DMV that have a `resource_associated_entity_id` of `72057594046971904`:

```
SELECT *
FROM sys.dm_tran_locks
WHERE resource_associated_entity_id = '72057594046971904';
```

On my test SQL Server, the query returned 17 locks, 15 of which are Exclusive locks on rows inside a clustered index; there is also one Intent Exclusive lock on a page and one Shared lock that is waiting to be placed on a page. The waiting Shared lock is the one returned by the `sys.dm_os_waiting_tasks` DMV. A portion of the results is displayed in Figure 7-7.

10	KEY	5	(291ba3e19696)	72057594046971904	0	X	LOCK	GRANT
11	PAGE	5	1:10672	72057594046971904	0	IX	LOCK	GRANT
12	PAGE	5	1:10672	72057594046971904	0	S	LOCK	WAIT
13	KEY	5	(6174a3e2309)	72057594046971904	0	X	LOCK	GRANT
14	KEY	5	(5851c0b34f7c)	72057594046971904	0	X	LOCK	GRANT

Figure 7-7. *Lock information from sys.dm_tran_locks*

Finding lock information and figuring out who is blocking whom by querying the `sys.dm_tran_locks` DMV can be a challenge on systems where you have many locks and blocks occurring, since the DMV will return a row for every lock placed. Another, easier, method to analyze locking and blocking is to use the `sp_whoIsActive` stored procedure created by Adam Machanic. This stored procedure will return information about everything that is running at that time and is a great tool for analyzing performance problems. With some extra parameters it will also return a wealth of locking information without you having to join various DMVs yourself. You can download the `sp_whoIsActive` DMV from Adam's blog at <http://sqlblog.com/files/default.aspx>.

To show you an example of the `sp_WhoIsActive` stored procedure I ran it on my test SQL Server while it was experiencing a blocking problem. The most basic way of running it is by just executing it:

```
EXEC sp_WhoIsActive;
```

Figure 7-8 shows a small portion of the results. There are many more columns available that will show you additional information, like the session ID of the blocking session.

	dd hh:mm:ss.mss	session_id	sql_text	login_name	wait_info
1	00 00:15:35.740	55	<?query -- BEGIN TRAN UPDATE Sales.SalesOrde...	EVDL-SQL2012-01\Administrator	NULL
2	00 00:15:28.906	56	<?query -- SELECT * FROM AdventureWorks2012 Sal...	EVDL-SQL2012-01\Administrator	(928759ms)LCK_M_S

Figure 7-8. `sp_WhoIsActive` default results

We can directly identify Wait Statistics information and the queries that are being executed at this time. Because the `blocking_session_id` column returned a session ID of 55 for the SELECT query, we should take a look at what the query that is being executed by session ID 55 is doing. By clicking on the query link inside the `sql_text` column we can view the whole query text, as shown in Figure 7-9.

```

1  | <?query --
2  |
3  | BEGIN TRAN
4  |
5  | UPDATE Sales.SalesOrderDetail
6  | SET CarrierTrackingNumber = '4E0A-4F89-AD'
7  | WHERE SalesOrderID = '43661'
8  | --?>

```

Figure 7-9. `sql_text` output from `sp_WhoIsActive`

In this specific case we can probably resolve the blocking problem pretty quickly. The query that is causing the block has left its transaction open without performing a COMMIT or ROLLBACK. As long as a transaction stays open, locks are being kept in place and are not released.

The `sp_WhoIsActive` stored procedure also has a parameter to return additional locking information, including information about the lock hierarchy. The query that follows will execute the `sp_WhoIsActive` stored procedure and retrieve extra lock information for the queries that are currently executing:

```
EXEC sp_WhoIsActive @get_locks=1;
```

Figure 7-10 returns the new columns that are added to the output of `sp_WhoIsActive`.

	locks	used_memory	status	open_tran_count	percent_complete
1	<Database name="AdventureWorks2012"><Locks><Loc...	2	sleeping	1	NULL
2	<Database name="AdventureWorks2012"><Locks><Loc...	3	suspended	0	NULL

Figure 7-10. Lock information returned by `sp_WhoIsActive`

By clicking the link below the `locks` column we can view which locks are being used and on what objects they are placed. This will also give us a good look at the locking hierarchy that is in place for this specific query. Figure 7-11 shows the extra lock information for first returned query in Figure 7-10.

```

179 <Object name="SalesOrderDetail" schema_name="Sales">
180 <locks>
181 <lock resource_type="KEY" index_name="Ida_CarrierTrackingNumber" request_mode="X" request_status="GRANT" request_count="30" />
182 <lock resource_type="KEY" index_name="PK_SalesOrderDetail_SalesOrderID_SalesOrderDetailID" request_mode="X" request_status="GRANT" request_count="15" />
183 <lock resource_type="OBJECT" request_mode="IX" request_status="GRANT" request_count="1" />
184 <lock resource_type="PAGE" page_type="***" index_name="Ida_CarrierTrackingNumber" request_mode="IX" request_status="GRANT" request_count="1" />
185 <lock resource_type="PAGE" page_type="***" index_name="PK_SalesOrderDetail_SalesOrderID_SalesOrderDetailID" request_mode="IX" request_status="GRANT" request_count="1" />
186 </locks>
187 </Object>
188 <Object name="SalesOrderHeader" schema_name="Sales">
189 <locks>
190 <lock resource_type="KEY" index_name="PK_SalesOrderHeader_SalesOrderID" request_mode="X" request_status="GRANT" request_count="1" />
191 <lock resource_type="OBJECT" request_mode="IX" request_status="GRANT" request_count="1" />
192 <lock resource_type="PAGE" page_type="***" index_name="PK_SalesOrderHeader_SalesOrderID" request_mode="IX" request_status="GRANT" request_count="1" />
193 </locks>
194 </Object>

```

Figure 7-11. Additional lock information returned by `sp_WhoIsActive`

I did scroll a bit down inside the lock information results to the interesting parts. Here we can see that an Intent Exclusive lock is placed at the OBJECT level, which means a lock has been placed on the object that is displayed in the `Object name` field of the XML schema—in this case `SalesOrderDetail`, which is a table. A level down, at the page level, we also see two Intent Exclusive locks in place inside two different indexes on the `SalesOrderDetail` table. And at the bottom level we see the Exclusive locks, which are placed at the KEY objects, which indicate rows inside an index.

There are many more parameters available for the `sp_WhoIsActive` stored procedure, each one of them returning more information about various parts of SQL Server. This makes the `sp_WhoIsActive` Stored Procedure a great tool for finding out what is going on inside your SQL Server instance, and I encourage you to give it a try.

Now that we have discussed many aspects of locking and blocking, from lock modes and hierarchies to analyzing locks and blocks, we should be ready to take a look at the lock-related Wait Types inside SQL Server. Keep in mind that this introduction to locking and blocking is far from a complete guide to the topic, as going into more detail on how locking and concurrency works inside SQL Server would fill a book by itself.

LCK_M_S

The first lock-related Wait Type is the `LCK_M_S` Wait Type. This Wait Type represents that a task is waiting to place a Shared lock on a resource.

What Is the LCK_M_S Wait Type?

The LCK_M_S Wait Type indicates that a task is, or has, been waiting to place a Shared lock on a resource. It is important to understand that you will only see this Wait Type when some form of blocking is occurring, since a task is waiting to place the Shared lock. It doesn't mean there is a Shared lock active on the resource. This is true for every lock-related Wait Type, as they will only get recorded when there is a blocking situation.

Since the LCK_M_S Wait Type is related to Shared locks, it will occur when a read action is being performed but has to wait because an incompatible lock is already in place on the resource we want to read. The time we are waiting before we are able to place the Shared lock is recorded as the Wait Time of the LCK_M_S Wait Type.

Figure 7-12 shows a common situation that will result in LCK_M_S waits occurring. In this case an Exclusive lock has been placed on a page by T1, indicating a data modification. When T2 wants to read the data from the page it will need to place a Shared lock, but since Exclusive and Shared locks are incompatible, a LCK_M_S wait occurs.



Figure 7-12. LCK_M_S wait occurring

LCK_M_S Example

Creating an example of a LCK_M_S wait occurring is not very difficult, as we just need to create a block situation between a data modification query and a data read query.

For this example we are going to run the query seen in Listing 7-1 against the AdventureWorks database. This query will begin a transaction and modify a few rows, but it will not commit or rollback the transaction. Since we explicitly indicated this transaction by supplying a `BEGIN TRAN`, SQL Server will keep the locks in place until we explicitly execute a `COMMIT` or `ROLLBACK` command.

Listing 7-1. Start a modification transaction

```
BEGIN TRAN

UPDATE Sales.SalesOrderDetail
SET CarrierTrackingNumber = '4EOA-4F89-AD'
WHERE SalesOrderID = '43661';
```

When we execute the query we receive a result very quickly; in my case 15 rows were updated. But like I said before, the transaction is not yet finished, so it will remain running, leaving locks on the objects it modified.

So far we aren't causing any blocking, since this is the test SQL Server and no other queries are running. Let's change that and create a blocking situation.

For this we are going to open a second window in SQL Server Management Studio and execute the query seen in Listing 7-2. This will just perform a SELECT against the `Sales.SalesOrderDetail` table, the same table in which we are currently modifying data.

Listing 7-2. Select data from a table where a modification is being performed

```
SELECT *
FROM AdventureWorks2012.Sales.SalesOrderDetail;
```

As soon as we run this SELECT query we notice that no results are returned and that the query will keep running. This is a typical example of a blocking operation where a transaction is modifying data we want to read inside another transaction.

If we were to query the `sys.dm_os_waiting_tasks` DMV we would be able to see the `LCK_M_S` Wait Type, as shown in Figure 7-13.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address	blocking_task_address	blocking_session_id
1	0x000000017673C928	55	0	25616	LCK_M_S	0x0000000178B154C0	NULL	54

Figure 7-13. `LCK_M_S` wait occurring

The only way the `LCK_M_S` wait will be resolved is if the incompatible lock is removed. In this case we performed a rollback of the modification transaction we started in the first SQL Server Management Studio window. We do this by running the `ROLLBACK` command in the same session window. Immediately after performing the transaction rollback we receive the results the SELECT query asked for. Querying the `sys.dm_os_waiting_tasks` also showed that the `LCK_M_S` wait was resolved.

Lowering `LCK_M_S` Waits

Seeing `LCK_M_S` waits occur does not necessarily have to mean something is wrong. It does, however, indicate that blocking is occurring. If you notice high Wait Times on the `LCK_M_S` Wait Type it means that someone's read transaction is currently taking a long time to complete because it has to wait to place the Shared lock. So the first step will be to identify the query that is causing the block. We can do this by using the `sys.dm_os_waiting_tasks` DMV and looking at the `blocking_session_id` column. This is relatively quick to do when there is only a single block active, but can get complex when many concurrent queries are being blocked by other transactions. In this case we have to follow the blocking chain until we find the head blocker (which is the first lock on an object). Another option is to use the `sp_WhoIsActive` stored procedure we discussed in the "Locking and Blocking Introduction" at the start of this chapter. This stored procedure will move through the blocking chain for you, directly displaying the head blocker.

After we have found the query that is causing the blocking to occur, we need to analyze it and see if we can optimize that query. Maybe it is requesting more locks than it actually needs and thus requires a long time to complete. One way to optimize that query would be to look at whether any indexes should be added so fewer rows are required to be locked. Or maybe you could cut the single transaction into multiple transactions that each access fewer objects. Another possible issue that can cause more locking than necessary is out-of-date statistics. Statistics are used as input for a query plan, and if they do not accurately reflect the contents of the table or index they can lead to a bad query plan, which in turn can lead to more locks than necessary.

Another option would be to change to Isolation level of the read transactions so no Shared locks are needed in order to read the data. For instance, setting the Isolation level to Read Uncommitted will not place Shared locks, and the read transaction will not be blocked. This does introduce another problem related to the Isolation level, dirty reads, which we discussed in the “Locking and Blocking Introduction” section of this chapter. Next to using Read Uncommitted, you could also use Snapshot Isolation, which will result in fewer Shared locks, but will not cause dirty reads. Snapshot Isolation does put more load on the TempDB database, since it must maintain versions of data if many concurrent transactions are modifying that data.

LCK_M_S Summary

The LCK_M_S Wait Type occurs when an incompatible lock is being placed on a resource and another transaction wants to place a Shared lock on the same resource. Seeing the LCK_M_S Wait Type means transactions are being blocked. You should try to identify which queries are causing the block to occur and see if these can be optimized to result in fewer locks or locks that have a shorter duration. As a final resort you could choose to change the Isolation level of your read transactions, though this does introduce other side effects, like dirty reads or increased load inside TempDB.

LCK_M_U

LCK_M_U Wait Types are related to locks that use the Update (U) mode. When a task wants to place an Update lock on a resource but an incompatible lock is already in place, LCK_M_U waits occur.

What Is the LCK_M_U Wait Type?

The Update lock type is a special type of lock mode that indicates that data modification is about to occur. Even though its name might suggest it is only related to UPDATE queries, Update locks can also appear when performing INSERT or DELETE statements.

Update locks primarily exist to prevent deadlocks from occurring. Deadlocks indicate that two transactions that want to modify the same object are waiting indefinitely on each other to acquire an Exclusive lock on the resource. To understand how a deadlock situation can occur, and how Update locks can prevent this, take a look at the following scenario that would occur when no Update locks are used.

When two concurrent transactions want to perform a modification on the same object, both transactions would first place a Shared lock on the resource while the data they intended to modify was located. Since Shared locks are compatible with other Shared locks, both transactions would not block each other. When one of the two transactions found the data it needed to modify, it would convert its Shared lock to an Exclusive lock, and then a problem would occur. Since Shared locks are incompatible with Exclusive locks, and since the other transaction would also have a Shared lock on the resource, the conversion from Shared lock to Exclusive lock would not occur. The transaction would need to wait until the Shared lock of the other transaction was removed before it could convert its own Shared lock to an Exclusive lock, but since the other transaction also wants to convert its Shared lock to an Exclusive lock, both transactions would end up waiting on each other, and a deadlock would occur. SQL Server will automatically detect deadlock situations and choose one of the deadlocked transactions as a victim and perform a rollback of that transaction, ending the deadlock situation. Figure 7-14 shows a graphical representation of that situation.

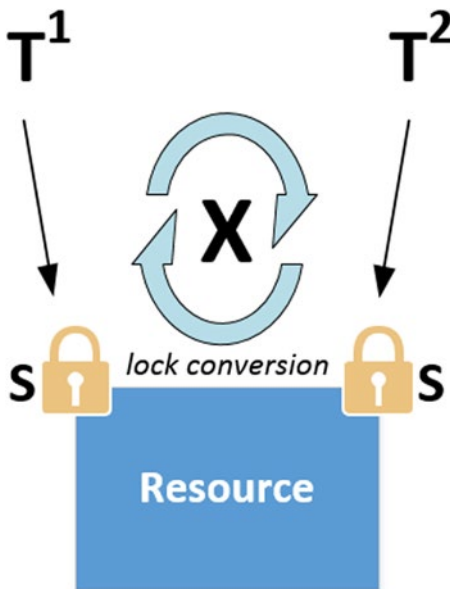


Figure 7-14. *Deadlock during lock conversion*

When Update locks are used inside SQL Server, no deadlock situation could occur. Update locks are compatible with Shared locks, but not with Exclusive or other Update locks. In the scenario above, the first transaction to find the data it needed to modify would not directly convert to an Exclusive lock, but rather would convert to an Update lock first. Since Update and Shared locks are compatible, there would be no problem converting to an Update lock, even though there was a Shared lock in place from the other transaction. The Update lock would then get converted to an Exclusive lock so the data modification could occur. Figure 7-15 shows this lock behavior.




Figure 7-15. Update locks during concurrent data modifications

When a transaction wants to place an Update lock but there is an incompatible lock already in place on the object, for instance, an Exclusive lock, the LCK_M_U Wait Type will be recorded.

LCK_M_U Example

To show you an example of LCK_M_U waits occurring, we have to create a situation where concurrent transactions want to modify the same resource. For this we are going to make use of the Ostress utility to execute an identical query using multiple connections. The query I am going to execute can be seen in Listing 7-3. This will perform an UPDATE against the Person.Address table inside the AdventureWorks database. I saved the query inside a .sql file named LCK_M_U.sql.

Listing 7-3. Modify the Person.Address table

```
UPDATE Person.Address
SET City = 'Los Angeles'
WHERE StateProvinceID = 9;
```

After saving the file I run the Ostress utility using the following command:

```
"C:\Program Files\Microsoft Corporation\RMLUtils\ostress.exe" -E
-dAdventureWorks2012 -i"C:\lck_m_u.sql" -n150 -r5 -q
```

This will create 150 concurrent connections, each one executing the query in Listing 7-3 five times. This should be enough to create some blocking.

While the Ostress utility is running I query the sys.dm_os_waiting_tasks DMV to find out what tasks are waiting. A small portion of the results are shown in Figure 7-16.

	walting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address	blocking_task_address	blocking_session_id
1	0x000000016BF52928	89	0	3616	LCK_M_U	0x000000016DDF4AC0	0x0000000168182CF8	74
2	0x000000017454BC38	134	0	2916	LCK_M_U	0x000000016A483A00	0x0000000168182CF8	74
3	0x000000016BF53498	128	0	2893	LCK_M_U	0x000000016A4ED9C0	0x0000000168182CF8	74
4	0x000000017F013C38	123	0	2794	LCK_M_U	0x000000016A3F7480	0x0000000168182CF8	74
5	0x000000017454B0C8	133	0	2749	LCK_M_U	0x000000016A409DC0	0x0000000168182CF8	74
6	0x000000016AD45868	146	0	2454	LCK_M_U	0x0000000168E2B1C0	0x0000000168182CF8	74

Figure 7-16. LCK_M_U waits occurring

As you can see in Figure 7-16, many different sessions are waiting to acquire an Update lock but are all being blocked by session ID 74. If we query the `sys.dm_tran_locks` DMV for lock information about this session, we can see it is granted an incompatible Exclusive lock, as shown in Figure 7-17.

request_mode	request_type	request_status	request_reference_count	request_lifetime	request_session_id
X	LOCK	GRANT	0	33554432	74
IX	LOCK	GRANT	0	33554432	178
S	LOCK	GRANT	1	0	57

Figure 7-17. Session ID 74 holding an Exclusive lock

All the other sessions will have to wait until the Exclusive lock of session ID 74 is removed. Then one of those sessions will acquire the Update lock it is requesting, convert it into an Exclusive lock, and perform its modification. That cycle will repeat until all the sessions are done with their modifications.

Lowering LCK_M_U Waits

Lowering LCK_M_U waits uses the same approach as lowering LCK_M_S Wait Types: try to identify the transaction that is causing the blocking to occur and try to optimize its locking behavior.

Changing the Isolation level will have little effect on LCK_M_U Wait Times since other Isolation levels have the most impact on transactions that perform reads. This makes optimizing your queries and/or indexes the way to go if you need to lower higher-than-normal Wait Times on the LCK_M_U Wait Type.

LCK_M_U Summary

The LCK_M_U Wait Type is related to locks that use the Update lock mode. Update locks are used to prevent deadlocks from occurring when concurrent transactions try to convert their Shared locks to Exclusive locks. Lowering LCK_M_U Wait Times is primarily achieved by optimizing potential blocking queries or indexes.

LCK_M_X

Another of the most common lock-related Wait Types is the LCK_M_X Wait Type. Just like both lock-related Wait Types we have already discussed, the LCK_M_X Wait Type is related to a specific lock type, in this case the Exclusive lock. And just like the other two lock-related Wait Types, seeing this Wait Type means there is some form of blocking occurring.

What Is the LCK_M_X Wait Type?

The LCK_M_X Wait Type occurs when a task is waiting to place an Exclusive lock on an object. Since Exclusive locks are not compatible with just about any other lock mode, including other Exclusive locks, seeing blocking occur when there are many concurrent modifications is pretty common. This means that seeing LCK_M_X waits occur is pretty common as well, especially in systems that have a high amount of concurrent transactions.

LCK_M_X Example

To demonstrate LCK_M_X waits occurring we are going to execute a SELECT statement without committing it. Before we run the SELECT we are going to set the Isolation level to Repeatable Read. Doing so makes sure the Shared locks are not removed while the transaction is still running. Since we do not end the transaction, the locks will remain on the objects until we either kill the transaction or perform a COMMIT or ROLLBACK. The query that follows shows the SELECT statement we will execute against the AdventureWorks database:

```
SET TRANSACTION ISOLATION LEVEL REPEATABLE READ
```

```
BEGIN TRANSACTION
```

```
SELECT *
FROM HumanResources.Employee;

-- COMMIT
```

Notice we commented out the COMMIT section to make sure the locks remain in place. Executing the query returns results pretty quickly, after just 1 second I got all the rows of the HumanResources.Employee returned. If we query the sys.dm_tran_locks DMV we should see that all the Shared locks are still in place, as shown in Figure 7-18.

	resource_type	resource_subtype	resource_database_id	resource_description	resource_associated_entity_id	resource_lock_partition	request_mode	request_type	request_status
10	KEY		5	(1bb936b3fc9)	72057594045136896	0	S	LOCK	GRANT
11	KEY		5	(d8b93f4a521)	72057594045136896	0	S	LOCK	GRANT
12	KEY		5	(6e3c257bc101)	72057594045136896	0	S	LOCK	GRANT
13	KEY		5	(ad3225e48be9)	72057594045136896	0	S	LOCK	GRANT
14	KEY		5	(#41092227a68)	72057594045136896	0	S	LOCK	GRANT
15	KEY		5	(#f79f0ee14f)	72057594045136896	0	S	LOCK	GRANT
16	KEY		5	(371e92bde080)	72057594045136896	0	S	LOCK	GRANT
17	KEY		5	(501004daffae)	72057594045136896	0	S	LOCK	GRANT

Figure 7-18. Shared locks still in place

While the locks are still in place we will run another query inside a new window in SQL Server Management Studio. The query that follows will perform an UPDATE on a single row inside the same `HumanResources.Employee` table:

```
UPDATE HumanResources.Employee
SET JobTitle = 'Tester'
WHERE BusinessEntityID = 5;
```

As soon as we execute the query above we'll notice a block occurring since the query keeps running without returning any results. This is as expected since there is a Shared lock in place on the row, or specifically the index key, that prevents us from updating it.

When we look at the `sys.dm_os_waiting_tasks` DMV, shown in Figure 7-19, we will notice that the query in the second window is waiting to place an Exclusive lock, indicated by the `LCK_M_X` Wait Type.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address	blocking_task_address	blocking_session_id
1	0x000000016A9E7C38	58	0	4583	LCK_M_X	0x0000000154386C80	NULL	54

Figure 7-19. *LCK_M_X wait occurring*

If we end the SELECT query, by either executing the COMMIT statement or by closing the windows inside SQL Server Management Studio, the Shared locks are removed and the second query will be able to execute its UPDATE command, ending the `LCK_M_X` wait.

Lowering LCK_M_X Waits

To lower `LCK_M_X` Wait Times you should use the same approach as for lowering other lock-related Wait Types. Try and identify what queries are causing the blocking and see if you can optimize them so they cause less blocking.

LCK_M_X Summary

The `LCK_M_X` Wait Type is related to Exclusive locks being blocked by other locks already in place on the same resource. Since Exclusive locks are incompatible with just about every other lock type, seeing `LCK_M_X` waits occurring is not uncommon for SQL Server instances that experience concurrent query execution.

LCK_M_I[xx]

Seeing the `LCK_M_I[xx]` Wait Type means that a task is being blocked when placing an Intent lock. Since we already discussed the various lock modes on objects, I replaced the lock mode used for the Intent lock as `[xx]` when discussing this Wait Type. The `[xx]` can be replaced by a variety of different lock modes; for instance, a block on an Intent Shared lock would be represented by the `LCK_M_IS` Wait Type, while a block on an Intent Exclusive lock would be shown as `LCK_M_IX`.

What Is the LCK_M_I[xx] Wait Type?

LCK_M_I[xx] Wait Types indicate that a task is waiting to place an Intent lock on an object. As we learned from the “Introduction to Locking & Blocking” section at the start of this chapter, Intent locks indicate that a lock of the same type is placed on an object lower down in the locking hierarchy. This doesn’t mean Intent locks are only there to warn SQL Server that there is a lock further down the hierarchy. Intent locks behave just like any other lock, and it is entirely possible that one Intent lock can block another, incompatible, Intent lock. Intent locks do have a little bit more flexibility regarding other incompatible Intent locks. For instance, it is possible for two Intent Exclusive locks to exist on the same page object, indicating that a row is going to be modified. It is even possible to have an Intent Shared lock on a page object together with an Intent Exclusive, because both of the locks can read and/or modify different rows.

Next to indicating the type of lock that exists lower down in the locking hierarchy, Intent locks have a few “special” modes the other lock modes do not have. It is possible for Intent locks to represent more than one lock mode on lower levels of the locking hierarchy. The list that follows describes these three Intent lock modes:

- Shared with Intent Exclusive (SIX). The Shared with Intent Exclusive lock mode represents that there are Shared Locks on all objects at a lower level, and Intent Exclusive locks on some of these objects. These locks are acquired by one transaction that wants to read data and plans to modify other data at the same time. When a task is being blocked while trying to place the SIX lock, it will be recorded by the LCK_M_SIX Wait Type.
- Shared Intent Update (SIU). This lock mode is a combination of Shared and Intent Update locks. Again, it is possible for a single transaction to acquire, and hold, both these lock modes at the same time at a lower level. If a block occurs while trying to place this lock, the LCK_M_SIU Wait Type will be used to record the Wait Time.
- Update Intent Exclusive (UIX). This lock mode is another combination of two other lock modes, Update and Intent Exclusive. Blocks on this lock mode will be represented by the LCK_M_UIX Wait Type.

Seeing high Wait Times on Intent locks is not very common, since Intent locks are a lot more flexible regarding their incompatibility with each other. This means there generally is less blocking on the Intent level than there is further down the locking hierarchy.

LCK_M_I[xx] Example

In this example we will generate a wait of the LCK_M_IX Wait Type. This means a transaction is waiting to acquire an Intent Exclusive lock on a higher level in the locking hierarchy.

We will use more or less the same example as we did for the LCK_M_X Wait Type by running a SELECT statement using the REPEATABLE READ Isolation level and not completing the transaction. The query that follows is the query I will be running against the AdventureWorks database, Person.Address table:

```
SET TRANSACTION ISOLATION LEVEL REPEATABLE READ

BEGIN TRAN

SELECT * FROM Person.Address;

--COMMIT
```

The COMMIT command has been commented out to leave the transaction open.

If we take a look at the sys.dm_tran_locks DMV, shown by Figure 7-20, we see that while the query is running there is only one lock currently active, a Shared lock on the OBJECT resource type. This indicates that the entire table is locked. Since this lock exists on this high level, there is no need for other Shared locks further down the hierarchy.

resource_type	resource_subtype	resource_database_id	resource_description	resource_associated_entity_id	resource_lock_partition	request_mode	request_type	request_status
6	DATABASE	5		0	0	S	LOCK	GRANT
7	DATABASE	5		0	0	S	LOCK	GRANT
8	OBJECT	5		373576369	0	S	LOCK	GRANT

Figure 7-20. Shared lock on a table

If another transaction wants to update a row inside the same table, it would first try to acquire an Intent Exclusive lock on the table and page level before it could acquire an Exclusive lock on the row level. The query that follows is such a transaction, and in this case we will try to update a single row:

```
UPDATE Person.Address
SET AddressLine1 = '1227 Shoe St.'
WHERE AddressID = 5;
```

You'll notice that the query above “hangs” as long as the SELECT query still has its Shared lock on the table. Even though Intent locks are in most cases compatible with other Intent locks on the same object, having, in this case, a Shared lock on the table level while trying to perform a data modification lower in the hierarchy will cause a block to occur. Shared locks and Intent Exclusive locks are not compatible.

If we look at the sys.dm_os_waiting_tasks DMV we should be able to see that the task to place the Intent Exclusive lock is waiting, as shown in Figure 7-21 below.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address	blocking_task_address	blocking_session_id
1	0x00000001767C2188	54	0	108373	LCK_M_IX	0x000000016D000040	NULL	56

Figure 7-21. LCK_M_IX wait occurring

Lowering LCK_M_I[xx] Waits

Just like with the other lock-related Wait Types we discussed earlier, try to focus on the queries that are causing the blocking when trying to lower LCK_M_I[xx] Wait Times. Because LCK_M_I[xx] waits only occur when incompatible locks are being held on objects higher in the locking hierarchy, it can be worth the time to investigate why those locks are placed so high in the hierarchy. Lock escalation can cause this to happen. Lock escalation occurs when it is more efficient for SQL Server to place a single lock higher in the locking hierarchy instead of locking many objects lower down. For instance, instead of placing thousands of Shared locks on rows, SQL Server can decide to place a single Shared lock on the table level. This requires far less resources to place and maintain than thousands of single locks. As a matter of fact, this is exactly what is occurring in the example I have shown you of the LCK_M_IX Wait Type. The Person.Address table we are querying with the SELECT query has more than 19,000 rows inside it. When we ran our SELECT * query against the table it would mean that at least 19,000 row locks would be needed. Because placing and holding that many locks would take a great deal of resources, SQL Server decided to place a single Shared lock on the table instead of 19,000 locks on the rows.

If we can rewrite the query so it requires fewer locks, for instance, by only selecting the first *x* rows instead of everything, SQL Server would probably choose to lock the rows again, instead of the entire table. The query that follows is the same one as we used in the example earlier, but in this case we are only selecting the top 2,000 rows:

```
SET TRANSACTION ISOLATION LEVEL REPEATABLE READ

BEGIN TRAN

SELECT TOP 2000 * FROM Person.Address;

--COMMIT
```

If we execute this query and look at the sys.dm_tran_locks DMV, we see the results shown in Figure 7-22.

	resource_type	resource_subtype	resource_database_id	resource_description	resource_associated_entity_id	resource_lock_partition	request_mode	request_type	request_status
13	KEY		5	(4c696a5d88d8)	72057594043826176	0	S	LOCK	GRANT
14	KEY		5	(d8b6f34a521)	72057594043826176	0	S	LOCK	GRANT
15	KEY		5	(109f3e47a4ee)	72057594043826176	0	S	LOCK	GRANT
16	KEY		5	(9c41e4137481)	72057594043826176	0	S	LOCK	GRANT
17	KEY		5	(54a823a0754e)	72057594043826176	0	S	LOCK	GRANT
18	KEY		5	(987d0e575a54)	72057594043826176	0	S	LOCK	GRANT
19	KEY		5	(5094c3e45b9b)	72057594043826176	0	S	LOCK	GRANT
20	KEY		5	(dc8a19b08f4)	72057594043826176	0	S	LOCK	GRANT
21	KEY		5	(1463d4038a3b)	72057594043826176	0	S	LOCK	GRANT

Figure 7-22. Shared locks on KEY objects

Instead of placing a Shared lock on the table, SQL Server decided to place a lock on every row, or in this case on every KEY object because of the clustered index. This means that other transactions can still access the table for data modifications as long as they do not modify the rows where a Shared lock is active.

LCK_M_I[xx] Summary

The LCK_M_I[xx] Wait Type is related to Intent locks, or rather, cases when another incompatible lock is blocking the placement of an Intent lock. Intent locks are placed on higher-level objects to indicate that a lock has been placed on a lower level in the locking hierarchy. Unlike the lock modes we discussed earlier that only represent one type of lock, Intent locks can represent different lock modes lower down in the locking hierarchy. One common cause of high Wait Times on LCK_M_I[xx] Wait Types is cases when SQL Server escalates lower-level locks to a higher-level lock. In this situation Intent locks will be blocked and cannot be acquired.

LCK_M_SCH_S & LCK_M_SCH_M

The last two lock-related Wait Types I want to discuss in this chapter are the LCK_M_SCH_S and the LCK_M_SCH_M Wait Types. Both of these Wait Types are related to locks that are being placed on tables, the so-called Schema locks. We didn't give a lot of attention to Schema locks earlier in this chapter, but since they can have a pretty big impact on Wait Times when they occur, I wanted to include them.

What Are the LCK_M_SCH_S & LCK_M_SCH_M Wait Types?

The LCK_M_SCH_S and LCK_M_SCH_M Wait Types are both related to Schema locks. Schema locks are placed at the table level to protect the table from modifications while queries access the table, or to prevent queries from accessing the table while it is being modified. There are two different types of Schema lock, Schema Stability (Sch-S) and Schema Modification (Sch-M). Each of them has a different Wait Type associated with them when a task is being blocked from placing a Schema Stability or Schema Modification lock. The LCK_M_SCH_S Wait Type is recorded when a Schema Stability lock has to wait before it can get placed, and the LCK_M_SCH_M Wait Type is recorded when a Schema Modification lock is waiting to get placed.

Both Schema locks have pretty extreme compatibility with other lock types. The Schema Stability lock is compatible with all other types of locks except for the Schema Modification lock. The Schema Modification lock, on the other hand, is incompatible with every other lock type, including Intent locks.

When using Schema Stability locks it is impossible to modify or change the table in any way while queries are currently reading or writing from or to that table. Because Schema Stability locks are compatible with every lock mode (except for Schema Modification), it is completely normal to see a Schema Stability lock on the table level together with, for example, an Intent Exclusive lock to indicate data modification is occurring on a lower level inside the table, as shown in Figure 7-23.

	resource_type	resource_subtype	resource_database_id	resource_description	resource_associated_entity_id	resource_lock_partition	request_mode	request_type	request_status
11	OBJECT		5		373576369	0	IX	LOCK	GRANT
12	OBJECT		5		373576369	0	Sch-S	LOCK	GRANT

Figure 7-23. Sch-S and IX lock on same table

Schema Modification locks are the opposite from Schema Stability locks, as they prevent any queries from accessing a table while a modification to the table is being performed.

LCK_M_SCH_S & LCK_M_SCH_M Example

For the first example, I am going to add a new column to an existing table, and just as we did in the examples earlier in this chapter, I am going to keep the transaction open by not supplying a COMMIT or ROLLBACK command. The query that follows adds an extra column to the Person.Address table in the AdventureWorks database, but I left the ROLLBACK command commented so the locks stay in place:

```
BEGIN TRAN

ALTER TABLE Person.Address
  ADD
  Test VARCHAR(10);

--ROLLBACK
```

In a new window in SQL Server Management Studio I am going to execute a simple SELECT query against the Person.Address table, like the one here:

```
SELECT *
FROM Person.Address;
```

If we take a look at the sys.dm_tran_locks DMV while both queries are running, we should be able to see if there is any blocking going on. Figure 7-24 shows a part of the output of a SELECT * query against the sys.dm_tran_locks DMV.

	resource_type	resource_subtype	resource_database_id	resource_description	resource_associated_entity_id	resource_lock_partition	request_mode	request_type	request_status
14	OBJECT		5		373576369	0	Sch-M	LOCK	GRANT
15	OBJECT		5		373576369	0	Sch-S	LOCK	WAIT

Figure 7-24. Sch-M and Sch-S locks

As you can see from Figure 7-24, the first query we started, with the goal of adding a column to the Person.Address table, resulted in a Sch-M lock on the table. The second SELECT query is waiting to receive a Sch-S lock on the same table.

If we query the sys.dm_os_waiting_tasks DMV we should see a task that is waiting on the LCK_M_SCH_S Wait Type. Figure 7-25 shows the output of sys.dm_os_waiting_tasks while both queries are running.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address	blocking_task_address	blocking_session_id
1	0x00000001767C2558	59	0	6451	LCK_M_SCH_S	0x000000016D004C40	NULL	56

Figure 7-25. LCK_M_SCH_S wait occurring

Just as we expected, the SELECT query is waiting to acquire its Schema Stability lock.

If we were to reverse the example by starting a read transaction and leaving it open, and then try to modify the same table, we should run into a LCK_M_SCH_M wait, since we can only acquire a Schema Modification lock when there are no active transactions inside the table we want to modify.

To show this situation I executed the query that follows. This starts a SELECT query with the Repeatable Read Isolation level, but I am leaving the transaction open so the locks stay in place:

```
SET TRANSACTION ISOLATION LEVEL REPEATABLE READ

BEGIN TRANSACTION

SELECT * FROM
Person.Address;

-- COMMIT
```

In a new window inside SQL Server Management Studio, I am going to execute the table modification query we used earlier to demonstrate the LCK_M_SCH_S Wait Type, but without leaving the transaction open:

```
ALTER TABLE Person.Address
ADD
Test VARCHAR(10);
```

As you will probably notice when executing the second query, nothing is returned and the query keeps running, a clear indication of a block occurring.

Let's take a look at the `sys.dm_tran_locks` DMV again to see what we can find out. Figure 7-26 shows the output on my test SQL Server.

	resource_type	resource_subtype	resource_database_id	resource_description	resource_associated_entity_id	resource_lock_partition	request_mode	request_type	request_status
8	OBJECT		5		373576369	0	S	LOCK	GRANT
9	OBJECT		5		373576369	0	Sch-M	LOCK	WAIT

Figure 7-26. Sch-M lock waiting to be acquired

In this case the table has a Shared lock on it from the SELECT query. Because we are selecting information from a pretty large table, SQL Server decided to place a table lock instead of placing locks on a lower level. Because a Schema Modification lock is incompatible with every other lock type, a block occurs, and we will have to wait until the Shared lock is gone before we can perform our table modification.

Looking at the `sys.dm_os_waiting_tasks` DMV shows us the results we are expecting, a LCK_M_SCH_M wait, as shown in Figure 7-27.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address	blocking_task_address	blocking_session_id
1	0x000000017f01d868	56	0	350150	LCK_M_SCH_M	0x000000016d004980	NULL	54

Figure 7-27. LCK_M_SCH_M wait occurring

Lowering LCK_M_SCH_S & LCK_M_SCH_M Waits

When you see waits occurring of either the LCK_M_SCH_S or LCK_M_SCH_M Wait Type, there is probably a transaction active that wants to modify the table. In the case of high Wait Times on the LCK_M_SCH_S Wait Type, the table modification transaction is already running; when seeing LCK_M_SCH_M waits, the modification is waiting for all active transactions to remove their locks on the table.

Modifying a table is not something that happens every day on production SQL Server instances (hopefully). Changing large tables can especially be problematic and a cause for high LCK_M_SCH_S Wait Times, and users that are trying to query the table that is being modified will notice delays. If, however, you absolutely need to modify a table, but there are some long-running queries retrieving information from that table, you can expect LCK_M_SCH_M waits.

Lowering the Wait Times of both Wait Types is directly related to performing modifications to tables. A suggestion could be to perform the table modification after office hours, or when there are as few as possible concurrent transactions accessing the table, instead of doing the modification when there are many transactions active against the table.

LCK_M_SCH_S & LCK_M_SCH_M Summary

The LCK_M_SCH_S and LCK_M_SCH_M Wait Types are the result of Schema Stability or Schema Modification locks being blocked by other locks. Seeing high Wait Times of either Wait Type indicates that either a table modification is waiting for all active locks on that table to be removed, or a table modification is currently running and other transactions are being blocked by it.

CHAPTER 8



Latch-Related Wait Types

In the previous chapter we took a pretty deep look at locking and blocking inside SQL Server, together with different Wait Types that indicate blocking is occurring. Latches look a lot like the locks we discussed earlier; in some cases they even appear to use the same modes as the locks we discussed in the previous chapter. Make no mistake though, latches are completely different than locks, even though they seem to share some features. While locks are used to guarantee transactions are isolated and consistent, latches are used to guarantee the consistency of in-memory objects.

Latches are, just like locking and blocking, a pretty complex subject inside SQL Server. Latches even have their own latch-statistics DMV that records how much time has been spent waiting on specific latch types.

Because of the complexity of latches and their function inside SQL Server, I believe they require an introduction to better understand how you can troubleshoot latch-related Wait Types later in this chapter. For this reason, we will start this chapter with an introduction to latches, just as we did with the introduction to locking inside Chapter 7, “Lock-Related Wait Types.”

Introduction to Latches

Microsoft describes latches as “lightweight synchronization objects that are used by various SQL Server components” on Books Online. This description is pretty vague, and there is a lot more depth to latches than the description would initially suggest.

The first thing that is important to understand about latches, which we lightly touched upon in the introduction of this chapter, is that latches are completely different than locks. I have heard and read various discussions about latches that treat latches as if they were locks. This confusion is easily explained, as latches, at first glance, do look similar to locks as regards their behavior and naming conventions within SQL Server. Just like locks, latches have various “modes,” and some of the acronyms to indicate the type of mode used are the same as for some lock modes. Another thing locks and latches have in common is that both objects play a role in keeping SQL Server objects consistent, and the way they manage this seems identical. While locks are used to make sure transactions are consistent, protecting the transaction for the entire duration it is running, latches

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-1-4842-1139-7_8](https://doi.org/10.1007/978-1-4842-1139-7_8)) contains supplementary material, which is available to authorized users.

are only used for the duration they are necessary and are not bound to the duration of a transactions. During the duration of one transactions, many different latches will be acquired and released again. Figure 8-1 visualizes this behavior.




Figure 8-1. Lock and latch behavior during transactions

Placing and especially maintaining locks on SQL Server objects is an expensive operation, mostly because they need to stay in place during entire transactions. Because latches are only needed for specific operations, and are then released again, they are far less costly to use than locks. This explains the “lightweight” part in the latch definition Microsoft uses.

The second part of the latch definition, “synchronization object,” we discussed earlier in this book, but not under that exact name. If you have read through this book so far, you should have noticed, especially in Chapter 5, “IO-Related Wait Types,” that SQL Server uses various methods to handle concurrent threads accessing objects. In Chapter 5 we talked about mutual exclusion, which makes sure only one thread at a time can access a memory object. In the same chapter we also discussed semaphores that implement gates to limit concurrent access to memory. Latches are another method used to make sure concurrent threads do not threaten the consistency of in-memory objects, and it does this in a way that looks a lot like locking.

Latch Modes

Latches have five different modes available to use when accessing objects. The list that follows describes these five modes, some of which might look familiar.

- SH. The SH mode represents a Shared latch. This mode is used when the latch is reading page data.
- UP. The UP latch mode is used by Update latches that are used whenever a page needs to be modified. By using the Update latch the page can still be read by other latches.
- EX. The EX latch mode, or Exclusive latch, is also used when page modification occurs. Unlike the Update latch, the Exclusive latch does not allow read or write access by other latch modes.

- **KP.** The KP latch mode is used by Keep latches. Keep latches are used to protect the page so it cannot be destroyed by the Destroy latch. They are compatible with every other latch mode except for the Destroy mode.
- **DT.** Indicates Destroy latches. Destroy latches are used when removing contents from memory; for instance, when SQL Server wants to free up a data page in memory.

As you can see in this list, the first three latch modes look a lot like those used by locks, and function more or less the same way. And just like lock modes, latch modes are compatible or incompatible with other latch modes. The table that follows shows the latch compatibility matrix and whether the different modes are compatible with each other or not.

Table 8-1. Latch Compatibility Matrix

	SH	UP	EX	KP	DT
SH	Yes	Yes	No	Yes	No
UP	Yes	No	No	Yes	No
EX	No	No	No	Yes	No
KP	Yes	Yes	Yes	Yes	No
DT	No	No	No	No	No

Unlike locks, which can partly be controlled by Isolation levels and query hints, latches are completely controlled by the SQL Server engine. This means we cannot modify latch behavior like we can for locks.

Latch Waits

Whenever a latch has to wait because its request couldn't be granted immediately, a latch wait occurs. These waits are tracked and recorded by SQL Server inside the `sys.dm_os_wait_stats` DMV, and also inside a dedicated DMV that records specific latch Wait Times, `sys.dm_os_latch_waits`, which we will discuss in more detail a bit further down in this chapter.

Figure 8-2 shows a situation in which a latch wait occurs. In this example we are waiting for a data page to be read from the storage subsystem into the buffer cache. In this case latches are used to make sure the same data page on the storage subsystem is not being read into the buffer cache by multiple threads. While the latch is waiting for the page to read into memory, the `PAGEIOLATCH_SH` Wait Type will be recorded.

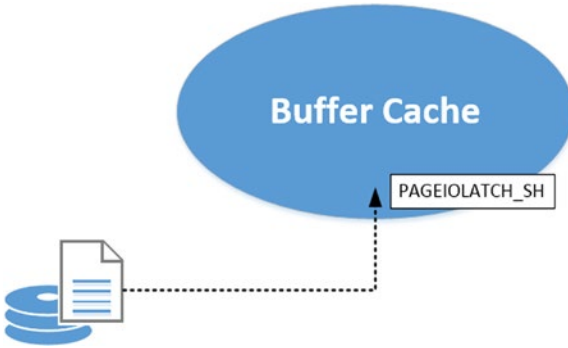


Figure 8-2. *PAGEIOLATCH_SH* occurring

There are three different latch Wait Types defined in SQL Server that can be accessed by querying the `sys.dm_os_wait_stats` DMV, and they are described in the list below.

- Buffer latches. Buffer latches are used to protect data pages inside the buffer cache. They are not only used for user-related data pages, but also for system pages like the Page Free Space (PFS) page that tracks free space inside data pages. Inside the `sys.dm_os_wait_stats` DMV they are indicated by the `PAGELATCH_[xx]` Wait Type, where the `[xx]` indicates the latch mode used.
- Non-buffer latches. These latches are used to protect data structures outside of the buffer cache. They are indicated by the `LATCH_[xx]` Wait Type inside the `sys.dm_os_wait_stats` DMV.
- IO latches. IO latches are used when data pages are read from the storage subsystem into the buffer cache. This type is indicated by the `PAGEIOLATCH_[xx]` Wait Type.

Figure 8-3 shows the different latch Wait Types recorded by the `sys.dm_os_wait_stats` DMV.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	LATCH_NL	0	0	0	0
2	LATCH_KP	0	0	0	0
3	LATCH_SH	5	0	0	0
4	LATCH_UP	0	0	0	0
5	LATCH_EX	669	318	6	82
6	LATCH_DT	0	0	0	0
7	PAGELATCH_NL	0	0	0	0
8	PAGELATCH_KP	0	0	0	0
9	PAGELATCH_SH	48	0	0	0
10	PAGELATCH_UP	0	0	0	0
11	PAGELATCH_EX	115	6	2	6
12	PAGELATCH_DT	0	0	0	0
13	PAGEIOLATCH_NL	0	0	0	0
14	PAGEIOLATCH_KP	0	0	0	0
15	PAGEIOLATCH_SH	1171	2699	173	79
16	PAGEIOLATCH_UP	25	14	1	0
17	PAGEIOLATCH_EX	639	468	165	5
18	PAGEIOLATCH_DT	0	0	0	0

Figure 8-3. Latch Wait Types inside `sys.dm_os_wait_stats`

Whenever you are looking at the LATCH_[xx] Wait Type inside the `sys.dm_os_wait_stats` DMV, you are actually looking at a summary of the Wait Times for these non-buffer latches. There are various non-buffer latch classes inside SQL Server, and to make it easier to analyze these non-buffer latch classes in more detail, the `sys.dm_os_latch_stats` DMV was added.

sys.dm_os_latch_stats

The `sys.dm_os_latch_stats` closely resembles the `sys.dm_os_wait_stats` DMV. The `sys.dm_os_latch_stats` DMV also shows the number of times a wait occurred, the total Wait Time, and the maximum Wait Time. The only column missing compared to the `sys.dm_os_wait_stats` DMV is `signal_wait_time_ms`; this is missing because latches do not follow the same execution process (RUNNING, SUSPENDED, RUNNABLE) as requests do.

Figure 8-4 shows a part of the `sys.dm_os_latch_waits` DMV. There are many more non-buffer latch classes, totaling 163, in SQL Server 2014.

	latch_class	waiting_requests_count	wait_time_ms	max_wait_time_ms
1	ACCESS_METHODS_DATASET_PARENT	670	320	7
2	ACCESS_METHODS_HOBT_FACTORY	0	0	0
3	ACCESS_METHODS_HOBT	3	1	1
4	ACCESS_METHODS_HOBT_COUNT	4	0	0
5	ACCESS_METHODS_HOBT_VIRTUAL_ROOT	0	0	0
6	ACCESS_METHODS_CACHE_ONLY_HOBT_ALLOC	0	0	0
7	ACCESS_METHODS_BULK_ALLOC	0	0	0

Figure 8-4. `sys.dm_os_latch_waits`

Just like the `sys.dm_os_wait_stats` DMV, the `sys.dm_os_latch_stats` DMV is cumulative since the start of the SQL Server service. This means it will get reset to 0 values again whenever your SQL Server service is restarted. We can also use the `BCC SQLPERF` command against the `sys.dm_os_latch_stats` DMV to reset the Wait Times manually by executing this command:

```
BCC SQLPERF('sys.dm_os_latch_stats', CLEAR)
```

Page-Latch Contention

One of the most common problems encountered regarding latches is page-latch contention. Page-latch contention occurs when many concurrent latches try to acquire a latch, but there already is a latch in place with an incompatible mode, causing a latch wait. Because this problem can occur on every SQL Server instance that is subjected to concurrent workloads, I want to provide you with the knowledge needed to identify page-latch contention before we discuss the various latch-related Wait Types.

There are a variety of things that can cause page-latch contention to occur, and even though we have little influence on the latch placement (remember, latches are placed and held by an internal process inside the SQL Server engine), the design of our database can impact latch behavior. One common cause for latch contention is when concurrent queries access so-called hot-spots inside your database. For instance, a small table that holds a few rows that need to be accessed by an application for configuration information can be a potential hot-spot. If many concurrent requests need data from this table, many latches will probably run into other, incompatible latches, causing latch waits to occur and slowing down the application's performance. I have seen this problem occurring various times for different clients, making this a real-world scenario, and I will show you an example of page-latch contention that is based on one of those cases.

In this case the client ran an application that, at specific times, would select large amounts of data and place the results into temporary tables. The application used a large amount of concurrent connections to speed up the creation of these temporary tables. To show the effects of this example, I am going to reproduce the scenario using Ostress to select rows from a table and then insert them into a temporary table.

As input for the Ostress utility I save a .sql file named `latch_contention.sql` with the query shown in Listing 8-1.

Listing 8-1. Select rows from `Sales.SalesOrderDetail` into temporary table

```
SELECT TOP (20000) *
INTO #tmptable
FROM Sales.SalesOrderDetail;
```

This query selects the top 20,000 rows from the `Sales.SalesOrderDetail` table inside the `AdventureWorks` database and inserts them into a temporary table (`#tmptable`).

The next step is to fire up Ostress and execute the `latch_contention.sql` script with 300 concurrent connections. The Ostress command line I use is shown here:

```
"C:\Program Files\Microsoft Corporation\RMLUtils\ostress.exe" -E
-dAdventureWorks -i"C:\latch_contention.sql" -n300 -r1 -q
```

While the Ostress utility is running, let's take a look at the `sys.dm_os_waiting_tasks` DMV to see if anything is running into waits. Figure 8-5 shows a part of the results that are returned when the query that follows is executed:

```
SELECT
    session_id,
    wait_duration_ms,
    wait_type,
    resource_description
FROM sys.dm_os_waiting_tasks;
```

	session_id	wait_duration_ms	wait_type	resource_description
10	88	2	PAGELATCH_UP	2:1:8088
11	91	37	PAGELATCH_UP	2:1:8088
12	82	56	PAGELATCH_UP	2:1:8088
13	93	41	PAGELATCH_UP	2:1:8088
14	95	5	PAGELATCH_UP	2:1:8088
15	98	89	PAGELATCH_UP	2:1:8088
16	99	53	PAGELATCH_UP	2:1:8088
17	100	58	PAGELATCH_UP	2:1:8088
18	103	33	PAGELATCH_UP	2:1:8088
19	117	1	PAGELATCH_UP	2:1:8088
20	116	53	PAGELATCH_UP	2:1:8088
21	105	20	PAGELATCH_UP	2:1:8088

Figure 8-5. *PAGELATCH_UP* waits occurring

We are running into a lot of `PAGELATCH_UP` waits here that indicate that a latch is waiting to update a page in-memory. The `resource_description` column is very useful here since it indicates the page ID that the latch wants to access. In this case the page ID we are trying to access is `2:1:8088`. The first number, 2, represents the database ID, which is the TempDB database. The second number, 1, indicates the file ID. Finally, the last number indicates the page ID, 8088. Why are all those sessions waiting on the same Wait Type against the same data page? This page happens to be a very special page, the Page Free Space (PFS) page. The PFS page tracks how much free space is left inside every page inside the database. It is always the first page of every database (page ID of 1) and has an interval of 8088 pages. So in this example, all the requests are waiting to update the second PFS page inside the TempDB database.

Because we are running inserts into a temporary table using many concurrent connections, we need to find, or allocate, data pages with free space to hold our rows inside the TempDB. All this space usage needs to be updated inside the PFS page, and latches are used to make sure only one thread gets access to the PFS page at a time. Figure 8-6 shows a Perfmon graph of two Perfmon counters, Transaction and Latch Waits/sec. This will show the relationship between to-Ostress workload and the number of latch waits occurring.




Figure 8-6. Latch Waits/sec and Transactions Perfmon graph

This is actually a classic example of the page-latch contention inside the TempDB database that can occur when many concurrent queries are creating objects inside TempDB. One way to resolve this specific case of latch contention is by adding more (equally sized) TempDB data files. Every new data file will maintain its own PFS pages, and adding more data files helps spread the load of updating the PFS pages. Using the query that follows, I added three more data files to the TempDB database:

```
USE [master]
GO

ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdev2', FILENAME =
N'D:\Data\tempdb2.mdf' , SIZE = 204800KB , FILEGROWTH = 10%)
GO
```

```
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdev3', FILENAME =
N'D:\Data\tempdb3.mdf' , SIZE = 204800KB , FILEGROWTH = 10%)
GO
```

```
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdev4', FILENAME =
N'D:\Data\tempdb4.mdf' , SIZE = 204800KB , FILEGROWTH = 10%)
GO
```

When running the Ostress utility again with the same query to insert data into a temporary table as we did before, I still notice latch waits occurring on PFS pages, but they are spread better across the TempDB data files. When looking at the Transactions and Latch Waits/sec Perfmon counters I also see fewer latch waits occurring, as shown in Figure 8-7.




Figure 8-7. Latch Waits/sec & Transaction Perfmon graph after adding more TempDB files

By adding more TempDB data files I would be able to lower the amount of latch waits even further. Adding too many TempDB data files can be a bad idea though, since the round-robin algorithm that makes sure the data files receive equal allocations can generate noticeable overhead when it needs to manage many TempDB data files. Paul Randal over at his SQLskills blog has a great post discussing TempDB data files and latch contention, which you can find here: <http://www.sqlskills.com/blogs/paul/a-sql-server-dba-myth-a-day-1230-tempdb-should-always-have-one-data-file-per-processor-core/>.

Now that we have discussed what latches are and how they work, and looked at an example of latch contention, let's move on and look at latch-related Wait Types.

PAGELATCH_[xx]

The first latch-related Wait Type we will discuss in this chapter is the `PAGELATCH_[xx]` Wait Type, where the `[xx]` indicates the latch mode used (for example, SH for Shared). Since we already discussed the various latch modes in the introduction of this chapter, we won't describe them again in this chapter.

What Is the `PAGELATCH_[xx]` Wait Type?

`PAGELATCH_[xx]` waits occur whenever a latch has to wait before it can access a page in-memory. The main cause for these waits are other latches that are already in place on the page and are incompatible with the latch mode our request wants to use. Just like a lock, the latch we want to place on the page has to wait until the incompatible latch is removed from the page. As long as the incompatible latch is in place, our request will record `PAGELATCH_[xx]` Wait Time. Figure 8-8 shows a graphical representation of a `PAGELATCH_UP` wait occurring. I used a cogwheel icon to indicate a latch is already in place on the page to avoid confusion with SQL Server locks.




Figure 8-8. `PAGELATCH_UP` wait occurring

It's easy to confuse `PAGELATCH_[xx]` waits with `PAGEIOLATCH_[xx]` waits. Even though they look alike in name, both are completely different latch Wait Types. The former indicates access to pages already in memory, the latter indicates pages are being read from disk into memory. We will go into detail regarding the `PAGEIOLATCH_[xx]` Wait Type a bit later in this chapter.

`PAGELATCH_[xx]` Example

In the introduction to this chapter we took a look at page-latch contention that can occur inside the TempDB database when many concurrent queries are loading data in temporary queries. This isn't the only form of latch contention that can occur inside SQL Server. One other form of latch contention is known as "last-page insert contention." Just like the page-latch contention scenario we discussed earlier, last-page insert contention can also be identified by noticing a high number of `PAGELATCH` waits. Let's go through an example of last-page insert contention.

Remember that time in database design class when you learned that every table should have a clustered index? And that the best candidate for a clustered index key column is a narrow, unique, ever-increasing value, like an integer? All of that is still true, and it absolutely helps to optimize the performance of queries against those tables. There are, however, very specific cases where using this practice can cause a performance problem known as last-page insert contention.

Last-page insert contention can occur on databases that experience a very heavy insert workload against a table with relatively small rows; for instance, a table with an ID column (Integer data type, auto increasing) and a Name column (Varchar data type). From a best-practice point of view, we would create a clustered index on the ID column since it fits the description of a good index key perfectly. It is narrow, unique for every row, and always increasing. But because of the ever-increasing nature of the auto increment, every newly added row will be added at the end of the clustered index, creating a hot-spot for the last data page of the clustered index. Figure 8-9 shows the insert behavior of rows into data pages inside a clustered index, inside the form of a so-called B-tree structure, which is the data structure SQL Server uses to sort indexes.

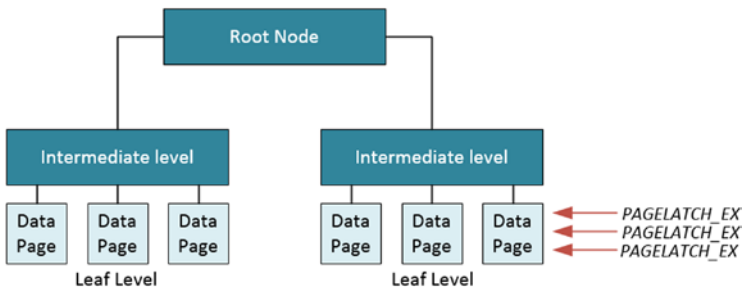


Figure 8-9. Last-page insert contention on last page in a clustered index

Even if the current data page is full, and a new data page is added, the target of the inserts will change to the new page, switching the hot-spot to the new data page.

One question I often hear about this behavior is: “Why aren’t locks stopping this?” The answer is actually pretty simple: because by default we will be using Exclusive row-level locks to insert our new rows instead of locking the page, and you can have multiple concurrent Exclusive row locks on one page. Access to the page that’s in-memory still needs to occur serially though, so latches are used to make sure only one thread has access to the page at any time. Figure 8-10 shows an enlarged view of the data pages at the leaf level of the clustered index with locks in place.

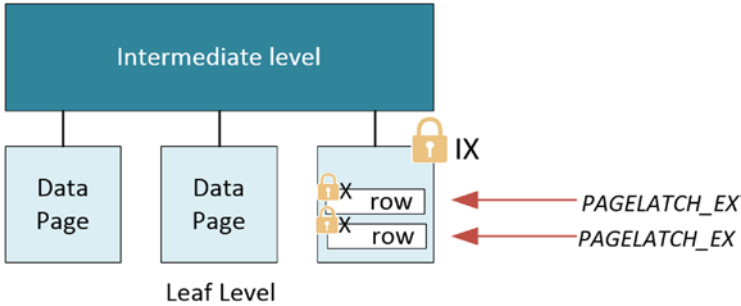


Figure 8-10. Leaf page of clustered index with locks in place

To show you an example of last-page insert contention, I will create a new table inside the AdventureWorks database of my test SQL Server instance using the query that follows:

```
CREATE TABLE Insert_Test
(
  ID INT IDENTITY (1,1) PRIMARY KEY,
  RandomData VARCHAR(50)
);
```

As you can see, this is a pretty small table with an ID column that automatically increases for every new row inserted, and a RandomData column that will hold some data. I indicate that the ID column is the primary key of this table, which will automatically create a clustered index using the ID column as the index key.

The next step is running Ostress with a highly concurrent workload that inserts new rows into the Insert_Test table. This time I don't create a .sql input file for Ostress but rather enter the following query in the Ostress command line:

```
INSERT INTO Insert_Test
(RandomData)
VALUES
(
  CONVERT(varchar(50), NEWID())
)
```

This will create the following Ostress command line that will connect to the AdventureWorks database and execute the query we supplied using 500 concurrent connections, each connection performing the query 100 times. This should create enough concurrent inserts to demonstrate last-page insert contention:

```
"C:\Program Files\Microsoft Corporation\RMLUtils\ostress.exe" -E
-dAdventureWorks -Q"INSERT INTO Insert_Test (RandomData) VALUES
(CONVERT(varchar(50), NEWID()))" -n500 -r100 -q
```


While Ostress is running, I query the `sys.dm_os_waiting_tasks` DMV:

```
SELECT
    session_id,
    wait_duration_ms,
    wait_type,
    resource_description
FROM sys.dm_os_waiting_tasks;
```

This query filters out some columns so that a screenshot of the results will fit on the page. Figure 8-11 shows a portion of the results.

	session_id	wait_duration_ms	wait_type	resource_description
9	279	41	PAGELATCH_EX	5:1:24342
10	73	42	PAGELATCH_EX	5:1:24342
11	336	42	PAGELATCH_EX	5:1:24342
12	323	43	PAGELATCH_EX	5:1:24342
13	250	43	PAGELATCH_EX	5:1:24342
14	490	44	PAGELATCH_EX	5:1:24342
15	500	44	PAGELATCH_EX	5:1:24342
16	496	44	PAGELATCH_EX	5:1:24342
17	497	45	PAGELATCH_EX	5:1:24342
18	551	45	PAGELATCH_EX	5:1:24342
19	547	45	PAGELATCH_EX	5:1:24342
20	546	45	PAGELATCH_EX	5:1:24342
21	544	45	PAGELATCH_EX	5:1:24342
22	128	44	PAGELATCH_EX	5:1:24342

Figure 8-11. *PAGELATCH_EX* waits on the same page

As expected, the insert workload caused a hot-spot to appear on a page inside the clustered index, in this case the page with a page ID of 24342. All of those tasks shown in Figure 8-11 (and there were around 300 more not shown) are all waiting to place an Exclusive page latch, indicated by the `PAGELATCH_EX` Wait Type, on that page so they can perform their insert operation.

To prove that page 24342 is a data page I am going to use the undocumented `DBCC IND` command to show us the pages that are associated with the `Insert_Test` table. `DBCC IND` will return a row for every page associated with the table we supply as a parameter to `DBCC IND`, and, among other things, will show us the page type of every page returned.

Running the command that follows will execute the `DBCC IND` command against the AdventureWorks database's `Insert_Test` table. Before we run the actual `DBCC IND` command we have to enable `Traceflag 3604` so the results of the `DBCC IND` command get returned in the SQL Server Management Studio results tab:

```
DBCC TRACEON (3604);
GO
```

```
DBCC IND (AdventureWorks2012, Insert_Test, 1);
GO
```

Figure 8-12 shows the results of the `DBCC IND` command. Highlighted is page 24342, the page that was the insert hot-spot during the Ostress workload.

	PageFID	PagePID	IAMFID	IAMPID	ObjectID	IndexID	PartitionNumber	PartitionID	iam_chain_type	PageType	IndexLevel
28	1	24338	1	24117	1198067454	1	1	72057594229489664	In-row data	1	0
29	1	24339	1	24117	1198067454	1	1	72057594229489664	In-row data	1	0
30	1	24340	1	24117	1198067454	1	1	72057594229489664	In-row data	1	0
31	1	24341	1	24117	1198067454	1	1	72057594229489664	In-row data	1	0
32	1	24342	1	24117	1198067454	1	1	72057594229489664	In-row data	1	0
33	1	24343	1	24117	1198067454	1	1	72057594229489664	In-row data	1	0
34	1	24344	1	24117	1198067454	1	1	72057594229489664	In-row data	1	0
35	1	24345	1	24117	1198067454	1	1	72057594229489664	In-row data	1	0
36	1	24346	1	24117	1198067454	1	1	72057594229489664	In-row data	1	0

Figure 8-12. *DBCC IND results*

The information we are interested in resides in the `IndexID` and `PageType` columns. The `IndexID` column returns the Index ID that this page is associated with. We only have one index on the `Insert_Test` table, and it has an ID of 1. The `PageType` column returns the page type of the specific page. In this case the `PageType` of page 24342 is 1, which indicates that the page is a data page.

Remember, `DBCC IND` is an undocumented SQL Server command, and I included it here to show you information about the page where the last-page insert contention was occurring. I strongly advise against using it on production servers.

Lowering `PAGELATCH_[xx]` Waits

So far I have showed you two examples where `PAGELATCH_[xx]` waits can occur, page-latch contention on the PFS page of the `TempDB` database and last-page insert contention. There is another latch-contention problem that can occur when inserting rows into a small table with an index. This case of latch contention can also be identified by `PAGELATCH_[xx]` waits occurring, but it also has a connection with the `LATCH_[xx]` Wait Type. For this reason I am saving the explanation and example of this specific case of latch contention for the next section of this chapter where we will discuss the `LATCH_[xx]` Wait Type.

Lowering `PAGELATCH_[xx]` waits can be challenging. Frequently, they are related to the design of your database or your workload, and these can prove difficult to change in production environments. There are, however, a number of factors that can contribute to latch contention that are worth taking the time to check.

It is more common to see latch contention occurring on systems that have a large number of logical processors (16+) and high concurrent OLTP workloads. However, having fewer logical processors does not mean latch contention cannot occur. The examples of latch contention I have shown you so far in this chapter have all been generated on a virtual machine with only four logical processors. I had to create a high enough concurrent workload to reach latch contention. Having more logical processors means there are more threads available to perform work, which also results in more concurrent latches being placed, increasing the chances of latch contention. Adding logical processors when experiencing latch contention can, in this specific case, cause even more latch contention to occur instead of resolving it. Lowering the number of logical processors isn't an option either, because this will slow down all your other workloads.

The best way to resolve latch contention is by identifying where the contention is occurring and what type of contention you are dealing with.

If you are dealing with PFS page contention, a good first step would be to check if you are using one or multiple database data files. If you are using one database data file, a first step would be to add additional, equally sized data files and measure if this lowers the amount of `PAGELATCH_[xx]` waits occurring. If you already have multiple database data files you could try adding more, but be careful not to add too many, because having this can introduce other performance problems. Your goal should be to find a database data file “sweet spot” where you have enough database data files to minimize the impact of latch contention, but not so many as to cause the overhead to become too high. This depends entirely on your workload, so it is impossible for me to give you a generalized recommendation.

When dealing with last-page insert contention you could consider changing the index key to something else instead of a sequentially increasing value, like a GUID. Using a GUID as an index key will result in a larger index because of the byte requirements of a GUID. Also, because GUIDs are entirely random, keeping the index in order requires more work than when dealing with an ever-increasing, sequential value. It can also have consequences for your applications or queries that possibly would need to be rewritten to accommodate the change in data type.

Other factors to consider that can impact latch contention are indexing strategies, page fullness, and the number of concurrent connections to the database. Also, identifying and optimizing the access patterns to the data inside the database can help immensely. For instance, if you know your workload consists of many very small inserts against a single table, it might be worth taking the time to see if you can combine some of the small inserts into a larger batch, effectively lowering the amount of latches needed.

One final option for resolving latch contention is using a method called hash partitioning. Hash partitioning splits up your table or index into various partitions based on a value that is generated by using a computed column. Partitioning is only available in Enterprise Edition, but it is a method that can minimize, or completely prevent, latch contention.

Hash partitioning works by cutting up tables or indexes into partitions, with each partition holding a set of the data. Partitioning is frequently used for archiving data from inside a table to another filegroup that resides on other (cheaper) storage, while the current data resides on fast storage. In the case of hash partitioning, we are going to calculate a value for every row inside the table using a computed column. Based on that value, we will move the row to a partition.

The great advantage of using this method of partitioning indexes is that every partition has its own index tree. So even though the insert statements will still occur on the last (right-most) page of the index, it will be spread across the partitions. If we look at Figure 8-13 we see a case of last-page insert contention, where concurrent queries are trying to insert rows into an index as we discussed in the example above.

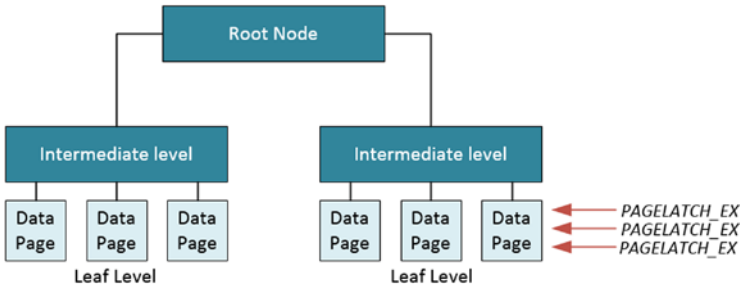


Figure 8-13. Last-page insert contention on the right-most data page of an index

If we were to use partitioning to cut the index into multiple parts (three in the image below), we would get the situation shown in Figure 8-14, multiple B-trees, each spanning a part of the data.




Figure 8-14. Last-page inserts spread across partitions

Let’s take a look at the effects of hash partitioning when we run the workload to generate last-page insert contention, like we did in the example earlier in this chapter.

The first thing we need to do is create a Partition function. This will map rows inside the table or index to partitions based on the value of a column. The script below will create a Partition function named LatchPartFunc that will divide rows into nine partitions based on the value of a column (which we will create a bit later). See the following:

```
CREATE PARTITION FUNCTION [LatchPartFunc] (INT)
AS RANGE LEFT FOR VALUES
(0,1,2,3,4,5,6,7,8);
```

The next step is to create a Partition scheme that will map the partitions to a filegroup:

```
CREATE PARTITION SCHEME [LatchPartSchema]
AS PARTITION [LatchPartFunc] ALL TO ([PRIMARY]);
```

In this case I used the PRIMARY filegroup, but you are free to create an additional filegroup to hold the partitions.

Next up is creating a new table called Insert_Test3 using the query that follows. Notice the ID_Hash column. This is a computed column that will calculate a value between 0 and 8 based on the value of the ID column:

```
CREATE TABLE Insert_Test3
(
  ID INT IDENTITY(1,1),
  RandomData VARCHAR(50),
  ID_Hash AS (CONVERT(INT, abs(binary_checksum(ID) % (9)), (0))) PERSISTED
);
```

The last step is to create a clustered index and map it to the Partition scheme:

```
CREATE UNIQUE CLUSTERED INDEX idx_ID
ON Insert_Test3
(
  ID ASC, ID_Hash
)
ON LatchPartSchema(ID_Hash);
```

Now that we have our partitioned table in place, let's repeat our Ostress workload that caused last-page insert contention in our previous example. I changed the target table for the inserts to our new, partitioned Insert_Test3 table.

```
"C:\Program Files\Microsoft Corporation\RMLUtils\ostress.exe"
-E -dAdventureWorks2012 -Q"INSERT INTO Insert_Test3 (RandomData)
VALUES (CONVERT(varchar(50), NEWID()))" -n500 -r100 -q
```

During both Ostress workloads I used Perfmon to monitor the number of latch waits occurring every second. Figure 8-15 shows the Perfmon graph for the first Ostress workload against a non-partitioned index and the second against the partitioned index we just created.

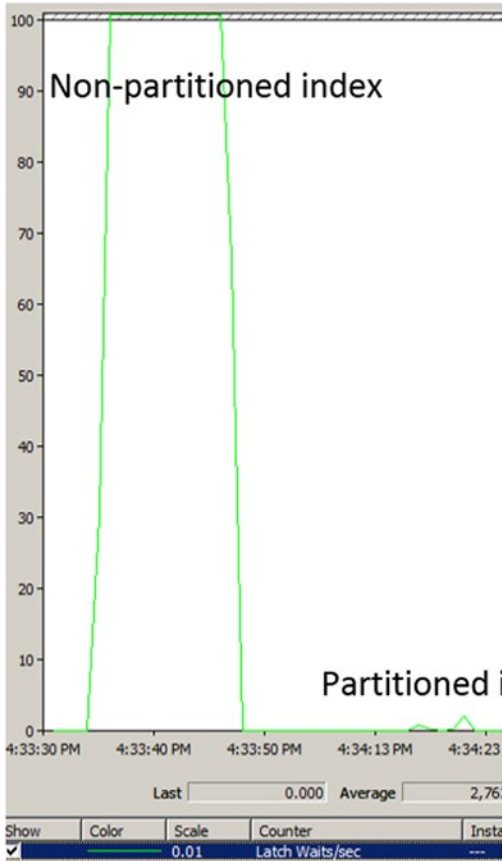


Figure 8-15. Latch Waits/sec against both a non-partitioned and a partitioned index

As you can see, the number of latch waits occurring dropped drastically after configuring hash partitioning! We can view the distribution of rows across the different partitions we created by running this query:

```
SELECT *
FROM sys.partitions
WHERE object_id = OBJECT_ID('Insert_Test3');
```

Figure 8-16 shows the results of this query on my test SQL Server instance.

	partition_id	object_id	index_id	partition_number	hobt_id	rows	filestream_filegroup_id	data_compression	data_compression_desc
1	72057594232111104	1598068879	1	1	72057594232111104	5555	0	0	NONE
2	72057594232178640	1598068879	1	2	72057594232178640	5556	0	0	NONE
3	72057594232242176	1598068879	1	3	72057594232242176	5556	0	0	NONE
4	72057594232307712	1598068879	1	4	72057594232307712	5556	0	0	NONE
5	72057594232373248	1598068879	1	5	72057594232373248	5556	0	0	NONE
6	72057594232438784	1598068879	1	6	72057594232438784	5556	0	0	NONE
7	72057594232504320	1598068879	1	7	72057594232504320	5555	0	0	NONE
8	72057594232569856	1598068879	1	8	72057594232569856	5555	0	0	NONE
9	72057594232635392	1598068879	1	9	72057594232635392	5555	0	0	NONE
10	72057594232700928	1598068879	1	10	72057594232700928	0	0	0	NONE

Figure 8-16. Rows distribution across partitions

In Figure 8-16 we see the nine partitions we created on the `Insert_Test3` table, numbered 1 to 9 by the `partition_number` column. The `rows` column shows the number of rows inside each partition, and as you can see, they are distributed very evenly across the nine partitions! The `hobt_id` returns the ID of the B-tree where two rows of this partition are stored; all the partitions have different IDs, meaning they each have their own B-tree structure.

Even though partitioning is a great way to resolve latch contention issues, it does come with its own unique challenges and drawbacks. Two of those are that it is an Enterprise-only feature, thus costly, and it can impact the generation of query execution plans, resulting in a sub-optimal plan.

PAGELATCH_[xx] Summary

The `PAGELATCH_[xx]` Wait Type indicates that buffer latches, which are used to protect in-memory pages, are running into other, non-compatible, buffer latches. Just like locks, latches have different modes they use when protecting pages, and not all of these are compatible with each other. Seeing a large amount of `PAGELATCH_[xx]` wait occurring can indicate a case of latch contention. Resolving latch contention can be challenging and frequently requires making changes to the database design or queries.

LATCH_[xx]

Another latch-related Wait Type is the `LATCH_[xx]` Wait Type. Just like the `PAGELATCH_[xx]` Wait Type we discussed in the previous section, `LATCH_[xx]` waits are related to a specific latch class. While the `PAGELATCH_[xx]` Wait Type is related to latches that protect data structures inside the buffer cache, the `LATCH_[xx]` Wait Type is related to latches that are used to protect data structures outside of the buffer cache (but still inside the SQL Server memory).

What Is the LATCH_[xx] Wait Type?

When you see the LATCH_[xx] Wait Type occurring a specific class of non-buffer latches is running into a wait. The LATCH_[xx] is actually a summary of the Wait Time of those different non-buffer latch classes and not a latch type of its own. All of the different non-buffer latch classes that add to the Wait Time shown by the LATCH_[xx] Wait Type are recorded inside their own DMV, `sys.dm_os_latch_stats`. There are many different latch classes that the LATCH_[xx] Wait Type represents, totaling 163 in SQL Server 2014. Figure 8-17 shows the memory area where LATCH_[xx] waits can occur.

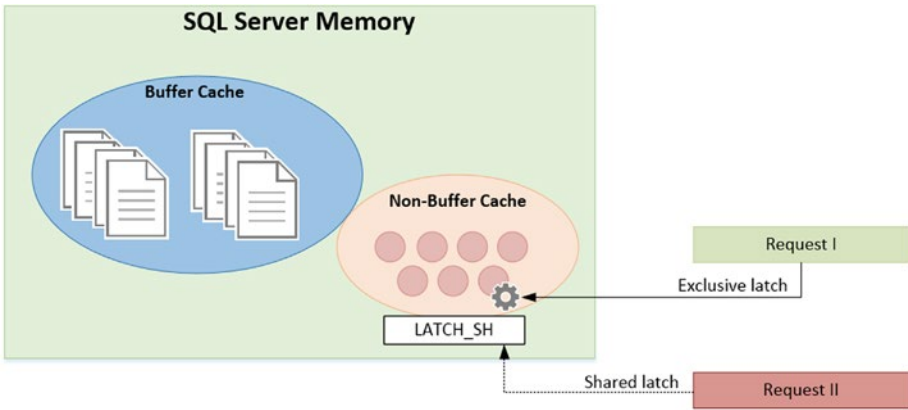


Figure 8-17. LATCH_SH wait occurring

Because the LATCH_[xx] Wait Type is a cumulative view of waits occurring on a specific latch class, you will need to look inside the `sys.dm_os_latch_stats` DMV to find the exact cause of the LATCH_[xx] wait. We described the inner workings and columns of the `sys.dm_os_latch_stats` DMV in the “Introduction to Latches” section at the start of this chapter, so I won’t go into more detail about the DMV here, but I included Figure 8-18 to show the underlying relationship between both DMVs.

wait_type	waiting_requests_count	wait_time_ms	latch_class	waiting_requests_count	wait_time_ms
1 LATCH_MG	0	0	ACCESS_METHODS_HOBT_VIRTUAL_ROOT	364168	17288047
2 LATCH_KP	0	0	BUFFER	301056	9449503
3 LATCH_SH	270760	11595603	LOG_MANAGER	1298	130879
4 LATCH_UP	0	0	ACCESS_METHODS_HOBT_COUNT	17	192
5 LATCH_EX	94672	5820874			
6 LATCH_DT	0	0			

Figure 8-18. Relationship between `sys.dm_os_wait_stats` and `sys.dm_os_latch_stats`

LATCH_[xx] Example

There is one case of latch contention that can occur that will result in LATCH_[xx] waits. This problem can occur on small tables that have a shallow B-tree structure (we will explain more about the B-tree structure a bit further down in this section) during a large volume of concurrent insert operations. A typical use case of such a table could be a messaging table that acts as a queue and gets truncated when the messages are sent. The script in Listing 8-2 will create a test table, named Insert_Test2, together with a non-clustered index on the table.

Listing 8-2. Test contention table with non-clustered index

```
-- Create the table
CREATE TABLE Insert_Test2
(
  ID UNIQUEIDENTIFIER,
  RandomData VARCHAR(50)
);

-- Create a non-clustered index on the ID column
CREATE NONCLUSTERED INDEX idx_ID
ON Insert_Test2 (ID);
GO
```

The ID column has a data type of UNIQUEIDENTIFIER to make sure random, non-sequential values are generated. By creating a non-clustered index on this column we are sure inserts will happen randomly across the B-tree associated with the non-clustered index.

Once the table is created we can start Ostress with a workload consisting of an insert query that will insert a single row inside the table. We will run the workload with 500 concurrent connections, with each of the connections executing the query 100 times. The command that follows shows the Ostress command:

```
"C:\Program Files\Microsoft Corporation\RMLUtils\ostress.exe" -E
-dAdventureWorks2012 -Q"INSERT INTO Insert_Test2 (ID, RandomData) VALUES
(NEWID(), CONVERT(varchar(50), NEWID()))" -n500 -r100 -q
```

While the workload is running I can take a look at the sys.dm_os_waiting_tasks DMV using the following query so the resource_description column could fit on the screenshot:

```
SELECT
  session_id,
  wait_duration_ms,
  wait_type,
  resource_description
FROM sys.dm_os_waiting_tasks;
```

Figure 8-19 shows a part of the results of this query on my test SQL Server instance.

	session_id	wait_duration_ms	wait_type	resource_description
7	254	4	LATCH_SH	ACCESS_METHODS_HOBT_VIRTUAL_ROOT (000000016BE6EF90)
8	88	4	LATCH_SH	ACCESS_METHODS_HOBT_VIRTUAL_ROOT (000000016BE6EF90)
9	170	3	LATCH_SH	ACCESS_METHODS_HOBT_VIRTUAL_ROOT (000000016BE6EF90)
10	160	3	LATCH_SH	ACCESS_METHODS_HOBT_VIRTUAL_ROOT (000000016BE6EF90)
11	283	3	LATCH_SH	ACCESS_METHODS_HOBT_VIRTUAL_ROOT (000000016BE6EF90)
12	432	1	LATCH_SH	ACCESS_METHODS_HOBT_VIRTUAL_ROOT (000000016BE6EF90)
13	312	3	LATCH_SH	ACCESS_METHODS_HOBT_VIRTUAL_ROOT (000000016BE6EF90)
14	202	2	LATCH_SH	ACCESS_METHODS_HOBT_VIRTUAL_ROOT (000000016BE6EF90)
15	113	1	LATCH_SH	ACCESS_METHODS_HOBT_VIRTUAL_ROOT (000000016BE6EF90)
16	281	2	LATCH_SH	ACCESS_METHODS_HOBT_VIRTUAL_ROOT (000000016BE6EF90)
17	251	1	LATCH_SH	ACCESS_METHODS_HOBT_VIRTUAL_ROOT (000000016BE6EF90)
18	296	1	LATCH_SH	ACCESS_METHODS_HOBT_VIRTUAL_ROOT (000000016BE6EF90)
19	365	0	LATCH_SH	ACCESS_METHODS_HOBT_VIRTUAL_ROOT (000000016BE6EF90)

Figure 8-19. *LATCH_SH* waits occurring

The `resource_description` column of the `sys.dm_os_waiting_tasks` DMV will help us identify what latch class is associated with the `LATCH_[xx]` wait. In this case we are running into the `ACCESS_METHODS_HOBT_VIRTUAL_ROOT` latch class.

Now that we know what latch class is running into waits, we can query the `sys.dm_os_latch_waits` DMV to find out the number of waits and the total Wait Time for this specific latch class using the following query:

```
SELECT *
FROM sys.dm_os_latch_stats
WHERE latch_class = 'ACCESS_METHODS_HOBT_VIRTUAL_ROOT';
```

Figure 8-20 shows the results of this query on my test SQL Server instance.

	latch_class	waiting_requests_count	wait_time_ms	max_wait_time_ms
1	ACCESS_METHODS_HOBT_VIRTUAL_ROOT	2134	61905	120

Figure 8-20. *ACCESS_METHOD_HOBT_VIRTUAL_ROOT* latch wait

Now that we have identified the latch class that is causing the `LATCH_SH` wait to occur we can start troubleshooting it. According to Books Online, the `ACCESS_METHODS_HOBT_VIRTUAL_ROOT` latch class is used to “synchronize access to the root page abstraction of an internal B-tree.” Even though the description is pretty limited it should give us an idea of where to start looking when troubleshooting this specific problem. I included a list of the different latch classes that are described on Books Online, including some extra information whenever possible, in Appendix III of this book.

Apparently something happened to the B-tree structures associated with indexes. Since we only have one non-clustered index on the table we created (`idx_ID`) we can assume something happened to the B-tree of that index. Let’s refresh our memory a little bit about what a B-tree index structure looks like by looking at Figure 8-21.

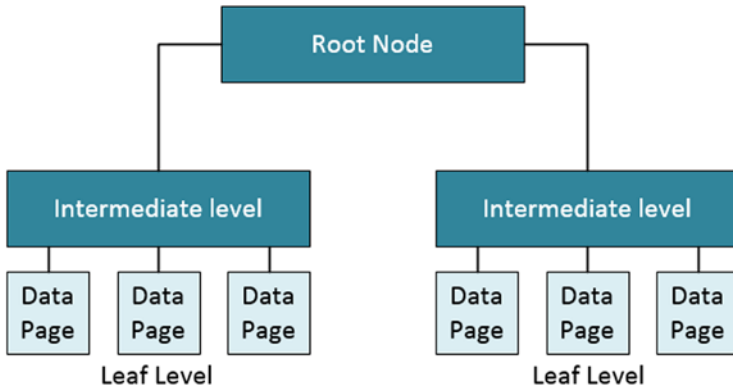


Figure 8-21. B-tree index structure

The B-tree structure we see in Figure 8-21 is pretty shallow, as it only has three levels. The first one is the Root Node (level 0), the second is the Intermediate level (level 1), and finally at the bottom of the B-tree are the data pages that hold the actual index keys (or in case of a clustered index, the entire row). In Figure 8-21 we only have one level of intermediate nodes, but depending on the number of data pages inside the index it is possible to have more Intermediate levels. When the table is very small it is possible to have the data pages inside the Intermediate level instead of a level further down the B-tree.

Whenever SQL Server needs to navigate through the B-tree it will start at the Root page inside the Root Node. The Root page will help it navigate down to the index page that holds the information it needs inside the Intermediate level. In turn, the Intermediate page can send the request further down the B-tree if the page is not inside the Intermediate level, but is rather at the Leaf level. The Leaf level is the last level in an index; it cannot navigate further down than that. Figure 8-22 shows how SQL Server navigates through the B-tree. In this case I used numbers as the index key to make it a bit easier.

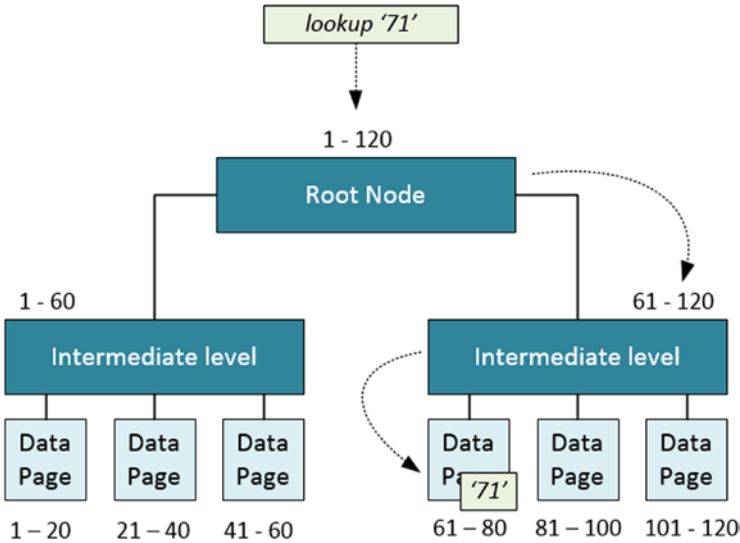


Figure 8-22. B-tree navigation

When data is added to the index, the index will allocate new data pages at the Leaf level to hold the index keys (remember, clustered indexes hold rows at the Leaf level). Whenever enough data is inserted that a new index level needs to be created, a Root page split occurs so the new level inside the index can be accessed. This Root page split doesn't cut your Root page into two new ones, there is always one Root page, but it needs to be updated so we can use it to navigate through the B-tree to the new data.

The ACCESS_METHODS_HOBT_VIRTUAL_ROOT latch class is the latch wait class that is associated whenever a Root page split occurs in order to create another level inside the B-tree. Whenever a Root page split occurs, the B-tree will acquire an Exclusive latch. All threads that want to navigate down the B-tree will have to wait for the Root page split to finish since they use Shared latches that are incompatible with the Exclusive latch. But why are we seeing LATCH_SH waits occurring when running our Ostress workload, instead of seeing Exclusive latches, since we are performing inserts? The reason for that is pretty simple: before SQL Server knows where to place the new index key inside the index, it first has to navigate through the B-tree to locate where the new index key needs to be placed, and it uses Shared latches during its navigation.

To show you that another level was added to the non-clustered index during our Ostress workload, I am going to use the INDEXPROPERTY function to retrieve the depth of the non-clustered index we created.

The first thing I am going to do is empty our Insert_Test2 table using the TRUNCATE command:

```
TRUNCATE TABLE Insert_Test2;
```

If we use the `INDEXPROPERTY` function against the non-clustered index on this table we can view the current depth of the B-tree. The query that follows shows how to use the `INDEXPROPERTY` function to retrieve this information:

```
SELECT INDEXPROPERTY(OBJECT_ID('Insert_Test2'), 'idx_ID', 'indexDepth')
```

Since we just truncated the table, the index depth should be 0 as there are no rows inside the table yet.

I then run the `Ostress` workload again, and after it has finished I look at the index information again. Instead of using the `INDEXPROPERTY` function, I use the `sys.dm_db_index_physical_stats` DMF to return some additional information about the number of index and data pages inside the index. This query returns such information:

```
SELECT
    index_id,
    index_type_desc,
    index_depth,
    index_level,
    page_count,
    record_count
FROM sys.dm_db_index_physical_stats
    (DB_ID(N'AdventureWorks'), OBJECT_ID(N'Insert_Test2'), NULL, NULL, 'DETAILED');
```

Figure 8-23 shows the results of this query on my test SQL Server instance.

	index_id	index_type_desc	index_depth	index_level	page_count	record_count
1	0	HEAP	1	0	420	50000
2	2	NONCLUSTERED INDEX	3	0	262	50000
3	2	NONCLUSTERED INDEX	3	1	2	262
4	2	NONCLUSTERED INDEX	3	2	1	2

Figure 8-23. `sys.dm_db_index_physical_stats` results

As you can see in this image, the non-clustered index now has three levels as indicated by the `index_depth` column. The `index_level` and `page_count` columns show how many pages exist on each level of the B-tree. The highest `index_level` number is the Root level, the lowest the Leaf level.

While the new levels were created inside the B-tree, the concurrent insert queries had to wait before they could navigate the B-tree, resulting in the `LATCH_SH` waits.

Lowering LATCH_[xx] Waits

In the previous example I presented a specific case of latch contention that occurs when index Root page splits occur so as to extend the B-tree structure. As I mentioned before, there are many, many more latch classes that are reported by the LATCH_[xx] Wait Type. This makes describing “one-size-fits-all” suggestions impossible. I can describe a general approach though, using the list here:

- Query `sys.dm_os_waiting_tasks` if LATCH_[xx] waits are occurring. The `resource_description` column can show you additional information about the specific latch class. If you are in a situation where the LATCH_[xx] waits do not show in `sys.dm_os_waiting_tasks` but high Wait Times are visible in `sys.dm_os_wait_stats` DMV, the `sys.dm_os_latch_waits` DMV should be your starting point. Another helpful DMV can be the `sys.dm_exec_requests` DMV. Joined together with the `sys.dm_exec_sql_text` DMF, it may help you to find the query that is causing the LATCH_[xx] wait.
- Query `sys.dm_os_latch_waits` to see if this correlates with the latch class shown in the `resource_description` column of the `sys.dm_os_waiting_tasks` DMV.
- Check Books Online or Appendix III in this book for more information about the specific latch class.

Another good resource for more information about common latch classes is Paul Randal’s blog post “Most common latch classes and what they mean” at <http://www.sqlskills.com/blogs/paul/most-common-latch-classes-and-what-they-mean/>. Though Paul only describes the ten most common latch classes, it can be a good starting point for your investigation.

Thankfully, it is not very common to see consistent high Wait Times for the LATCH_[xx] Wait Type since the cases that can cause the LATCH_[xx] waits to occur are frequently related to very specific workloads and database design.

LATCH_[xx] Summary

The LATCH_[xx] Wait Type represents waits encountered by a large selection of different, non-buffer-related latch classes inside SQL Server. These non-buffer-related latch classes have their own latch wait DMV, `sys.dm_os_latch_waits`, that returns the Wait Times of those latch classes. Troubleshooting LATCH_[xx] waits can be difficult since the latch classes that are associated with the Wait Type are minimally documented. Thankfully, it is not very common to see high Wait Times on the LATCH_[xx] Wait Type since they only occur for very specific situations and workloads.

PAGEIOLATCH_[xx]

The final latch-related Wait Type we will discuss in this chapter is the PAGEIOLATCH_[xx] Wait Type. The PAGEIOLATCH_[xx] Wait Type is by far the most common latch-related Wait Type and together with the CXPACKET Wait Type is the most common Wait Type to see on any SQL Server instance.

Just like the two previous latch Wait Types we discussed, the PAGEIOLATCH_[xx] has different access modes that I replaced with [xx] in this chapter. Since we already described the different latch modes in the introduction, we won't discuss them further in this chapter.

So far we have discussed two of the three latch-related Wait Types and the areas they are related to. The PAGELATCH_[xx] Wait Type was related to latches being placed on memory pages inside the buffer cache, and the LATCH_[xx] Wait Type is related to latches on non-buffer objects. The PAGEIOLATCH_[xx] Wait Type also indicates the use of latches on a specific area in SQL Server, in this case the IO latches.

What Is the PAGEIOLATCH_[xx] Wait Type?

Disk operations inside SQL Server are very expensive. Accessing the disk subsystem of your system requires extra resources and is always slower than accessing information that is inside the memory of your system. Because SQL Server is a database, and accessing and storing data inside the database is its primary function, the way SQL Server accesses data is extremely important. If data access is slow, SQL Server will perform slower as well, and this can result in noticeable performance degradation inside your queries or applications. To make IO interactions as efficient as possible, SQL Server uses a buffer cache to cache data pages that were previously accessed into the memory of your system. By caching data pages SQL Server only has to access the disk subsystem once, when the first query requests those specific data pages. When later queries require the same data pages as the first query, SQL Server will detect that those pages are already inside the buffer cache, through the Buffer Manager, and will access the data pages from inside the buffer cache instead of performing extra interactions with the disk subsystem. Figure 8-24 shows the buffer cache behavior when a query requires data pages from the storage subsystem.




Figure 8-24. Moving a page from the storage subsystem to the buffer cache

During the movement of data pages from the storage subsystem to the buffer cache latches are used to “reserve” a buffer page for the data page on the storage subsystem. This makes sure no other concurrent transactions allocate the same buffer page, or simultaneously attempt to transfer the same data page from the storage subsystem to the buffer cache.

While SQL Server is transferring the data page from the storage subsystem into the buffer cache an Exclusive latch will be placed on the buffer page. Because Exclusive latches are incompatible with almost every other latch mode (save for the Keep mode) it is guaranteed that no other latch can access the buffer page while it is being transferred. From the user perspective, a `PAGEIOLATCH_[xx]` wait will be recorded for the duration of the transfer of the data page. The mode of the latch depends on the action that initiated the movement of the data page from the storage subsystem to the buffer cache. A `PAGEIOLATCH_SH` will be recorded if the data is being moved for read access, and a `PAGEIOLATCH_UP` or `PAGEIOLATCH_EX` will be used if the data page is being moved for a modification. Figure 8-25 shows the data-page movement including latches and latch waits.

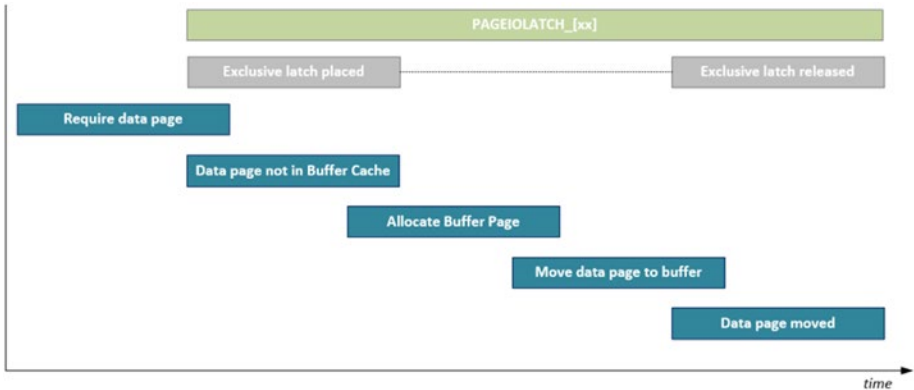


Figure 8-25. Movement of data page

To summarize the section above, if you see `PAGEIOLATCH_[xx]` waits occurring it means your SQL Server instance is reading data from your storage subsystem into your buffer cache. Because this is a very common operation to perform, it is easy to see why the `PAGEIOLATCH_[xx]` Wait Type is one of the most common Wait Types on any SQL Server instance.

PAGEIOLATCH_[xx] Example

Creating an example for `PAGEIOLATCH_[xx]` waits is extremely easy—just run a `SELECT` query against a freshly restarted SQL Server instance. A restart of the SQL Server service will empty the buffer cache of all data pages. This will leave you with a buffer cache without any user data inside it. There is, however, another way to clear the buffer cache without needing to restart the SQL Server service. Running the `DBCC DROPCLEANBUFFERS` command will remove all the unmodified data pages from the buffer cache. Combining it with the `CHECKPOINT` command will ensure the modified pages are also written to disk, leaving you with an empty, or “cold,” buffer cache.

The query in Listing 8-3 will perform a `CHECKPOINT`, followed by a `DBCC DROPCLEANBUFFERS`. It will then reset the `sys.dm_os_wait_stats` DMV and run a query against the AdventureWorks database. After the query against some of the tables inside the AdventureWorks database, we will query the `sys.dm_os_wait_stats` DMV for `PAGEIOLATCH_[xx]` waits.

Listing 8-3. Generate `PAGEIOLATCH_SH` waits

```
CHECKPOINT 1;
GO

DBCC DROPCLEANBUFFERS;
GO

DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR);
GO
```

```

SELECT
    SOD.SalesOrderID,
    SOD.CarrierTrackingNumber,
    SOH.CustomerID,
    C.AccountNumber,
    SOH.OrderDate,
    SOH.DueDate
FROM Sales.SalesOrderDetail SOD
INNER JOIN Sales.SalesOrderHeader SOH
ON SOD.SalesOrderID = SOH.SalesOrderID
INNER JOIN Sales.Customer C
ON SOH.CustomerID = C.CustomerID;

SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type LIKE 'PAGEIOLATCH_%';

```

Figure 8-26 shows the results of the last query against the `sys.dm_os_wait_stats` DMV on my test SQL Server instance.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	PAGEIOLATCH_NL	0	0	0	0
2	PAGEIOLATCH_KP	0	0	0	0
3	PAGEIOLATCH_SH	23	24	3	0
4	PAGEIOLATCH_UP	0	0	0	0
5	PAGEIOLATCH_EX	0	0	0	0
6	PAGEIOLATCH_DT	0	0	0	0

Figure 8-26. *PAGEIOLATCH_SH* Wait Time information

It makes sense to see only `PAGEIOLATCH_SH` waits recorded, since we are executing a `SELECT` query instead of performing data modification. If we were to perform some form of data modification against the rows we would see `PAGEIOLATCH_EX` waits in the results.

Lowering `PAGEIOLATCH_[xx]` Waits

If you read the section above you would probably understand that seeing `PAGEIOLATCH_[xx]` waits occur is completely normal behavior inside SQL Server. In many cases your database(s) are larger in size than the available amount of RAM inside your system, and some interaction with the storage subsystem is to be expected. Even if your databases are smaller than the amount of RAM in your system, and they can fit entirely inside the buffer cache, you will still notice `PAGEIOLATCH_[xx]` waits occurring during the startup of SQL Server, since this is the time SQL Server will start moving data pages from the storage subsystem into the buffer cache (if there is any query activity, SQL Server won't move data from the storage subsystem into the buffer cache by itself).

Since seeing PAGEIOLATCH_[xx] waits occur is completely normal for every SQL Server instance, it is very important to maintain a baseline of the Wait Times (Chapter 3, “Building a Solid Baseline,” can help you with that). When the Wait Times stay within range of the baseline values for this Wait Type, there shouldn’t be any cause for concern. If Wait Times are much higher than you expect them to be, investigation into the source of the higher-than-normal Wait Times might be necessary. There are quite a few possible causes for seeing higher-than-normal PAGEIOLATCH_[xx] Wait Times, and I will describe some of the more common ones.

The first place I look when noticing higher-than-normal PAGEIOLATCH_[xx] Wait Times is the SQL Server log to find out if SQL Server was restarted. SQL Server can restart due to a crash, but also when a failover occurs. These events will cause high PAGEIOLATCH_[xx] Wait Times that might not be reflected in your baseline, especially when SQL Server restarts do not frequently occur. Since our baseline’s measurements are frequently calculated using average values, the PAGEIOLATCH_[xx] Wait Times during SQL Server startup slowly lower when more measurements are taken inside the average baseline. If you create your baseline on measurements taken between a specific time range, and SQL Server hasn’t had a restart during the time range, your baseline measurements will also be considerably lower. As we read in the example section, a DBCC DROPLEANBUFFERS will also remove data pages from the buffer cache, resulting in higher PAGEIOLATCH_[xx] Wait Times after the command completes. Sadly, unlike the DBCC FREEPROCCACHE command, the execution of the DBCC DROPLEANBUFFERS command is not recorded in the SQL Server log.

One of the more common pieces of advice I see about lowering PAGEIOLATCH_[xx] Wait Times is to focus your attention on the storage subsystem. Since the PAGEIOLATCH_[xx] Wait Type indicates data movement from your storage subsystem to your buffer cache, it is logical that the storage subsystem plays a vital role in the Wait Times, but do not automatically assume this is the root cause! If you do have storage-related problems, this can show in the PAGEIOLATCH_[xx] Wait Time, so checking the performance of your storage subsystem is worth the effort.

A good place to start for monitoring storage performance is Perfmon. Perfmon has a variety of counters that will show you the current performance of your storage subsystem. Those in the list below are ones I use the most when monitoring storage performance.

- PhysicalDisk\Avg. Disk sec/Read. This will show you the average read latency on the disk you are monitoring. Less latency is better, and as a general guideline latency values should be below 20ms (0.020 within Perfmon, as it reports the latency in seconds).
- PhysicalDisk\Avg. Disk sec/Write. This will return the average write latency on the disk you are monitoring. Just like the read latency, write latency should, as a general guideline, be below 20ms.
- PhysicalDisk\Disk Reads/sec. Shows the amount of read IOPS (Input Output Operations) per second. This information can be helpful if you are running into capacity issues on the disk.

- PhysicalDisk\Disk Writes/sec. The same as the PhysicalDisk\Disk Reads/sec, but this one shows the amount of write IOPS.
- PhysicalDisk\Disk Read Bytes/sec. This counter shows the amount of bytes read from the disk per second. Again, this information can be useful for detecting possible capacity problems.
- PhysicalDisk\Disk Write Bytes/sec. This is identical to the PhysicalDisk\Disk Read Bytes/sec, but this counter shows the amount of bytes written to disk per second.

Using the information these Perfmon measurements provide, you should be able to identify possible storage-related bottlenecks. This information can also be helpful to the storage administrator (if there is one) who can compare these measurements to the measurements of the storage he/she manages.

As an extra diagnostic tool, or if you cannot use Perfmon, you can also run the IO performance script Paul Randal created based on the `sys.dm_io_virtual_file_stats` DMF, shown in Listing 8-4. The script, and the blog post describing the script, can be found on Paul's blog at <http://www.sqlskills.com/blogs/paul/how-to-examine-io-subsystem-latencies-from-within-sql-server/>.

Listing 8-4. IO performance script

```
SELECT
  [ReadLatency] =
    CASE WHEN [num_of_reads] = 0
      THEN 0 ELSE ([io_stall_read_ms] / [num_of_reads]) END,
  [WriteLatency] =
    CASE WHEN [num_of_writes] = 0
      THEN 0 ELSE ([io_stall_write_ms] / [num_of_writes]) END,
  [Latency] =
    CASE WHEN ([num_of_reads] = 0 AND [num_of_writes] = 0)
      THEN 0 ELSE ([io_stall] / ([num_of_reads] + [num_of_writes])) END,
  [AvgBPerRead] =
    CASE WHEN [num_of_reads] = 0
      THEN 0 ELSE ([num_of_bytes_read] / [num_of_reads]) END,
  [AvgBPerWrite] =
    CASE WHEN [num_of_writes] = 0
      THEN 0 ELSE ([num_of_bytes_written] / [num_of_writes]) END,
  [AvgBPerTransfer] =
    CASE WHEN ([num_of_reads] = 0 AND [num_of_writes] = 0)
      THEN 0 ELSE
        (([num_of_bytes_read] + [num_of_bytes_written]) /
         ([num_of_reads] + [num_of_writes])) END,
  LEFT ([mf].[physical_name], 2) AS [Drive],
  DB_NAME ([vfs].[database_id]) AS [DB],
  [mf].[physical_name]
```

```

FROM
    sys.dm_io_virtual_file_stats (NULL,NULL) AS [vfs]
JOIN sys.master_files AS [mf]
    ON [vfs].[database_id] = [mf].[database_id]
    AND [vfs].[file_id] = [mf].[file_id]
-- WHERE [vfs].[file_id] = 2 -- log files
ORDER BY [Latency] DESC
-- ORDER BY [ReadLatency] DESC
-- ORDER BY [WriteLatency] DESC;
GO

```

Figure 8-27 shows a part of the results of the IO performance script from Listing 8-4 on my test SQL Server instance, ordered by latency.

	ReadLatency	WriteLatency	Latency	AvgBPerRead	AvgBPerWrite	AvgBPerTransfer	Drive	DB	physical_name
1	19	2	19	64187	8192	63730	C:	master	C:\Program Files\Microsoft SQL Server\MSSQL11.MSS...
2	15	3	13	58380	43695	55941	D:	tempdb	D:\Log\templog.ldf
3	5	3	4	12572	99732	91114	D:	tempdb	D:\Data\tempdb2.mdf
4	6	3	4	14915	99642	90975	D:	tempdb	D:\Data\tempdb3.mdf
5	7	4	4	15425	99518	95162	D:	tempdb	D:\Data\tempdb4.mdf
6	5	4	4	17503	98285	94395	D:	tempdb	D:\Data\tempdb5.mdf
7	4	4	4	16533	102422	98121	D:	tempdb	D:\Data\tempdb5.mdf
8	6	4	4	22932	98442	91408	D:	tempdb	D:\Data\tempdb.mdf
9	4	3	3	104718	110074	106262	D:	AdventureWorks2012	D:\Data\AdventureWorks2012_Data.mdf
10	3	0	1	412333	5811	19045	D:	AdventureWorks2012	D:\Log\AdventureWorks2012_log.ldf

Figure 8-27. IO performance on my test SQL Server instance

One important thing to keep in mind with the IO performance script is that its values are cumulative from the start of the SQL Server service. They do not show the situation at the moment of executing the query. If you are interested in monitoring your IO performance using this script, you can consider capturing the output to a table at a specific interval and calculate the deltas (much in the same way as we did in Chapter 3, “Building a Solid Baseline”).

Next to the performance of your storage subsystem, the behavior of your queries can impact the Wait Times of the PAGEIOLATCH_[xx] Wait Type. The more data your queries are requesting (that is not already in the buffer cache), the larger the amount of data that needs to be read from the storage subsystem into the buffer cache. For instance, if we were to modify the query of Listing 8-3, shown in Listing 8-5, to be more selective so that less data is returned, the amount of PAGEIOLATCH_[xx] Wait Time should also be less.

Listing 8-5. Modified Listing 8-3 query

```

CHECKPOINT 1;
GO

DBCC DROPCLEANBUFFERS;
GO

DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR);
GO

```

```

SELECT
    SOD.SalesOrderID,
    SOD.CarrierTrackingNumber,
    SOH.CustomerID,
    C.AccountNumber,
    SOH.OrderDate,
    SOH.DueDate
FROM Sales.SalesOrderDetail SOD
INNER JOIN Sales.SalesOrderHeader SOH
ON SOD.SalesOrderID = SOH.SalesOrderID
INNER JOIN Sales.Customer C
ON SOH.CustomerID = C.CustomerID
WHERE SOD.CarrierTrackingNumber BETWEEN 'F467-41BF-8B' AND 'F4E4-4739-B4'
;

SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type LIKE 'PAGEIOLATCH_%';

```

Figure 8-28 shows the results of the last query that was executed against the `sys.dm_os_wait_stats` DMV.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	PAGEIOLATCH_NL	0	0	0	0
2	PAGEIOLATCH_KP	0	0	0	0
3	PAGEIOLATCH_SH	6	3	0	0
4	PAGEIOLATCH_UP	0	0	0	0
5	PAGEIOLATCH_EX	0	0	0	0
6	PAGEIOLATCH_DT	0	0	0	0

Figure 8-28. PAGEIOLATCH_SH Wait Time information

As you can see from Figure 8-28, the Wait Times of the PAGEIOLATCH_SH Wait Type went down drastically, from 24 milliseconds to 3 milliseconds. Now, this example is rather small, and we are dealing with very small result sets, but I think it shows the point.

However, you don't always have the luxury of being able to modify every query so that it is more selective. Maybe the queries are generated by an application and you can't even modify them, or the queries simply need the large result set. Thankfully, as a DBA, we can also play a part in minimizing PAGEIOLATCH_[xx] Wait Times by simply performing database maintenance. Index fragmentation and out-of-date statistics can increase the PAGEIOLATCH_[xx] Wait Times drastically. If indexes are fragmented more disk IOs need to take place to retrieve the data requested, which means IO latches will need to stay in place longer, which results in higher PAGEIOLATCH_[xx] Wait Times. Out-of-date statistics can also result in more disk IOs, because SQL Server expects a different amount of rows to be returned instead of the actual amount of rows. So, make sure you are regularly performing index and statistics maintenance to make sure the amount of disk interaction is as small as possible.

The final area that can impact PAGEIOLATCH_[xx] Wait Time is the memory of your system. SQL Server will remove data pages from inside the buffer cache if they have not been accessed within a specific timeframe in order to free up room inside the buffer cache. The interval at which SQL Server performs this cleanup depends on the amount of data coming into the buffer cache and the amount of free space inside the buffer cache. If the request for data pages inside the buffer cache is very high, SQL Server will be forced to swap data pages that have been accessed the least (or haven't been accessed for a while) for pages that are required now. This movement of data pages from and to the buffer cache will result in more PAGEIOLATCH_[xx] waits. In an ideal world, your database would fit completely inside the buffer cache of your SQL Server instance. In this case, SQL Server will only need to move the data pages from the storage subsystem into the buffer cache once, where they will stay until SQL Server restarts again. Even though we do have access to very large amounts of RAM these days, in many cases we cannot simply fit our entire database into the buffer cache of our SQL Server instance, and some swapping of data pages from the buffer cache back to the storage subsystem can be expected. Adding more RAM to your system will increase the number of data pages the buffer cache can store and can help the buffer cache keep those pages in memory longer.

There are two Perfmon counters that can help you get some insight into the buffer cache usage: SQLServer:Buffer Manager\Buffer cache hit ratio and SQLServer:Buffer Manager\Page life expectancy. The SQLServer:Buffer Manager\Buffer cache hit ratio will show you what percentage of pages could be located in the buffer cache that do not require a physical read on the storage subsystem. The SQLServer:Buffer Manager\Page life expectancy counter will show you the number of seconds a data page stays inside the buffer cache. If you see continuously low values on both these counters, compared to your baseline, it could mean SQL Server is running into memory pressure and needs to move data pages from the buffer cache back to disk again to free up memory. These two counters are not perfect, though, and much has been written about their workings (and specifically their ideal values). We won't go into details about what good values for these counters should be, as for that you should refer to your baseline, but I believe they are a good starting point for investigating buffer cache memory pressure.

PAGEIOLATCH_[xx] Summary

The PAGEIOLATCH_[xx] Wait Type is, by far, the most common latch-related Wait Type. Together with the CXPACKET Wait Type, the PAGEIOLATCH_[xx] Wait Type is probably the most common Wait Type on any SQL Server instance. The PAGEIOLATCH_[xx] Wait Type is directly related to the movement of data pages on the storage subsystem into the buffer cache memory of your SQL Server instance. SQL Server uses the buffer cache to minimize the number of interactions to the (much slower) storage subsystem so as to maximize performance. Whenever a data page is read into the buffer cache, the PAGEIOLATCH_[xx] Wait Type will be recorded for the time it took to do so. There are many methods available to lower the amount of PAGEIOLATCH_[xx] Wait Time. The frequently advised “get faster storage” doesn't always hold true, even though fast storage will indeed directly influence the PAGEIOLATCH_[xx] Wait Times. Optimizing queries so they require fewer data pages to be moved to the buffer cache, performing maintenance on indexes and statistics, and analyzing memory performance could all lead to lower PAGEIOLATCH_[xx] Wait Times.



High-Availability and Disaster-Recovery Wait Types

There have always been several options available within SQL Server to make sure your database is always available to your users and/or the data inside your database is replicated to another server so as to minimize the chances of losing data. Just like with performing regular database backups to ensure you can revert to a previous state of your database should a crash or data corruption occur, planning and maintaining highly-available database environments is part of your job as a DBA.

Now that data has become incredibly important for many companies, and the need for high-availability database servers grows, many DBAs will find themselves managing SQL Server instances inside a high-availability solution, like mirroring, or a disaster-recovery configuration, like log shipping. With these types of SQL Server high-availability and disaster-recovery configurations comes a group of dedicated Wait Types that are directly related to the health of your high-availability and disaster-recovery (HA/DR) configuration. With the release of SQL Server AlwaysOn Availability Groups in SQL Server 2012, more options became available for configuring HA/DR solutions, together with new Wait Types that are directly related to AlwaysOn Availability Groups.

In this chapter we will take a look at some of the most common Wait Types to see in HA/DR configurations. The main focus of the Wait Types inside this chapter is AlwaysOn Availability Groups, because Microsoft is deprecating many of the previous installments of features that now fall under the name AlwaysOn Availability Groups, like mirroring. As an exception to this rule, I selected one mirroring-related Wait Type that is relatively common on highly used mirroring configurations. All the other Wait Types are related to AlwaysOn Availability Groups.

For the examples inside this chapter, I used several virtual machines to create a mirroring and an AlwaysOn Availability Groups configuration. The configuration of these VMs can be found in Appendix I Example SQL Server Machine Configurations.

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-1-4842-1139-7_9](https://doi.org/10.1007/978-1-4842-1139-7_9)) contains supplementary material, which is available to authorized users.

DBMIRROR_SEND

The first Wait Type I want to discuss in this chapter is the DBMIRROR_SEND Wait Type. As you might suspect from the Wait Type name, DBMIRROR_SEND is related to database mirroring.

Database mirroring is a feature that was introduced in SQL Server 2005 but was announced deprecated in SQL Server 2012. This doesn't mean you cannot use database mirroring in SQL Server 2012 or SQL Server 2014, but it does mean it is scheduled for removal. The entire feature will be replaced with AlwaysOn Availability Groups, which offers the same configuration options as database mirroring.

Database mirroring is a solution that increases the availability of SQL Server databases, and unlike, for instance, failover clustering, it can be configured on a per-databases basis. Database mirroring works by redoing every data modification operation that occurs on the primary (called *principal* in database-mirroring terms) on the mirror database. The redoing of every database modification operation is achieved by streaming active transaction log records to the mirror server, which will perform the operations on the mirror database in the sequence in which they were inserted into the transaction log.

Database mirroring offers two different operating modes that impact the availability and performance of the mirror configuration. Synchronous (or high-safety) mode and Asynchronous (or high-performance) mode. Even though both modes perform identical actions to ensure data modification operations are also performed on the mirror database, there can be a large difference in performance, and thus in waits occurring.

The synchronous mirror mode makes sure that every data modification action that is performed on the principle is also directly performed on the mirror. It does this by waiting on sending a transaction confirmation message to the client until the transaction is successfully written to disk on the mirror. Figure 9-1 depicts synchronous mirroring.

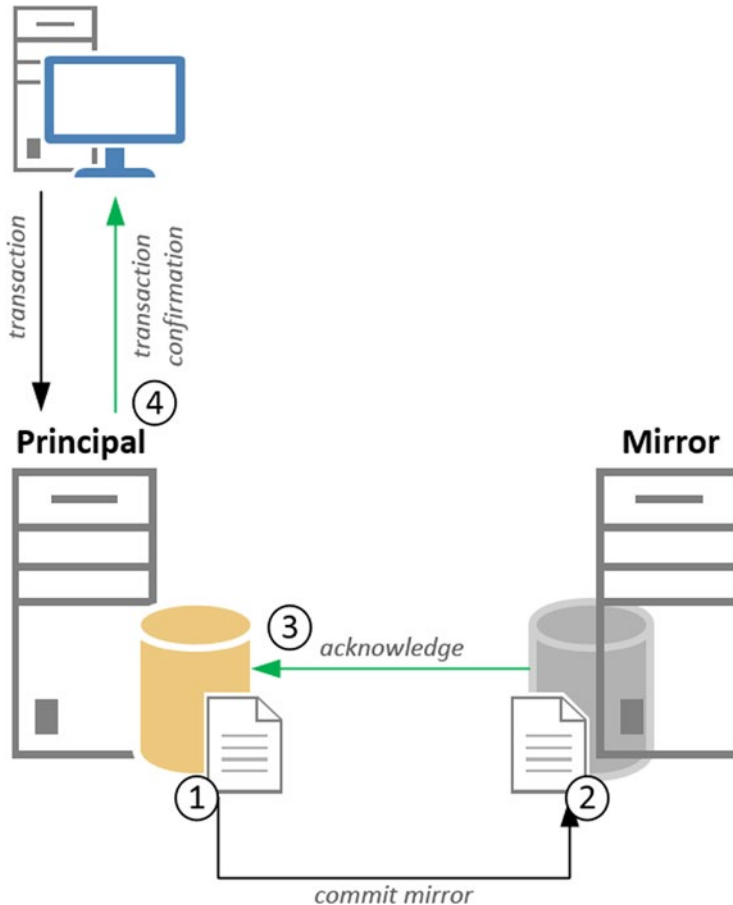


Figure 9-1. Synchronous mirroring

Even though synchronous mirroring makes sure data on the principal and mirror are 100% identical, it comes with a few drawbacks. One of those is that the performance of your database inside a synchronous mirroring configuration is highly dependent on the speed the mirror can process data modification operations, since every transaction has to be committed on the mirror first.

The flow of a data modification transaction is described in the steps that follow:

1. When the transaction is received, the principal will write the transaction to the transaction log, but the transaction is not yet committed though.
2. The principal will send the log record to the mirror.

3. The mirror will harden the log record to disk and send an acknowledgment to the principal.
4. After the principal receives the acknowledgment, it will send a confirmation message to the client that the transaction was completed.

The asynchronous mode works in much the same way; the exception is that it will not wait on an acknowledgment message from the mirror before sending the transaction confirmation message to the client. This means that transactions are committed to disk on the principal before they are written to disk on the mirror. Using asynchronous mirroring will improve mirror performance, since the latency overhead of synchronous mirroring is removed. The trade-off for this increase in performance is that asynchronous replication can lead to data loss in the case of a disaster, since it is possible that transactions were not yet committed on the mirror. Figure 9-2 shows the transaction-log flow on an asynchronous mirror; the dotted lines indicate that the actions are not performed directly.

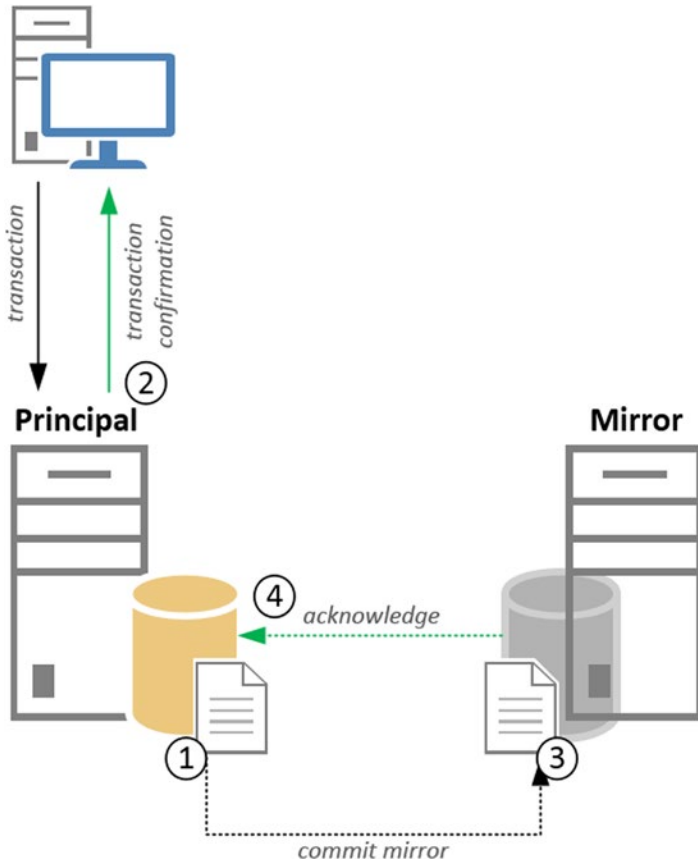


Figure 9-2. Asynchronous mirroring

What Is the DBMIRROR_SEND Wait Type?

The DBMIRROR_SEND Wait Type is most frequently related to synchronous mirroring configurations. The description of the DBMIRROR_SEND Wait Type on Books Online is “Occurs when a task is waiting for a communications backlog at the network layer to clear to be able to send messages. Indicates that the communications layer is starting to become overloaded and affect the database mirroring data throughput.” In this case the Books Online description is pretty accurate, but the network is not the only thing that can impact DBMIRROR_SEND Wait Times. Having a slow disk subsystem connected to the mirror database can also lead to an increase in DBMIRROR_SEND Wait Times.

Another important point to remember is that high DBMIRROR_SEND Wait Times will frequently only be recorded on the mirror instance, and not on the principal. It is common to see waits occur on the DBMIRROR_SEND Wait Type on both the principal and the mirror, but these will normally be very low on the principal. They can still reach high values on the mirror since, generally, there is always some latency between both SQL Server instances. Because of expected latency, I advise you to use baseline measurements to identify higher-than-normal Wait Times for the DBMIRROR_SEND Wait Type.

DBMIRROR_SEND Example

For this example I have build a synchronous mirror between two of my test SQL Server instances, using the AdventureWorks database as the database that will be mirrored between both instances.

Inside the AdventureWorks database I create a simple table using the script in Listing 9-1.

Listing 9-1. Create Mirror_Test table

```
USE [AdventureWorks]
GO

CREATE TABLE Mirror_Test
(
    ID UNIQUEIDENTIFIER PRIMARY KEY,
    RandomData VARCHAR(50)
);
```

After the table is created, I clear the `sys.dm_os_wait_stats` DMV and insert 10,000 rows into the `Mirror_Test` table using the query in Listing 9-2. I also make sure to clear the `sys.dm_os_wait_stats` DMV on the mirror as well before running the script in Listing 9-2.

Listing 9-2. Insert 10,000 rows into Mirror_Test table

```
DBCC SQLPERF('sys.dm_os_wait_stats, CLEAR')

INSERT INTO Mirror_Test
(
  ID,
  RandomData
)
VALUES
(
  NEWID(),
  CONVERT(VARCHAR(50), NEWID())
);
GO 10000
```

While the script is running I look at the DBMIRROR_SEND Wait Times on both the mirror and the principal using the query here:

```
SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'DBMIRROR_SEND'
```

The results of this query can be seen in Figure 9-3, which shows the DBMIRROR_SEND Wait Times on the mirror. Figure 9-4 shows the DBMIRROR_SEND Wait Times on the principal server.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	DBMIRROR_SEND	19176	32030	8057	1228

Figure 9-3. DBMIRROR_SEND Wait Times on mirror

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	DBMIRROR_SEND	0	0	0	0

Figure 9-4. DBMIRROR_SEND Wait Times on Principal

As you can see, we spend quite some time waiting on the DBMIRROR_SEND Wait Type on the mirror versus no DBMIRROR_SEND waits on the principal.

Lowering DBMIRROR_SEND Waits

One of the most common pieces of advice for lowering DBMIRROR_SEND Wait Time is changing the mirror mode from synchronous to asynchronous. While this will absolutely lower the Wait Time, it also means you can potentially lose data when a disaster occurs on the principal. Lowering the Wait Time on the DBMIRROR_SEND Wait Type will have a positive effect on the duration of your queries. For instance, in the example in the previous section the insert of 10,000 rows took around 40 seconds on my test SQL Server mirror configuration. When I changed the mirror mode from synchronous to asynchronous, not only did the Wait Times on the DBMIRROR_SEND Wait Type go down, the total execution time of 10,000 inserts went down to 10 seconds. That's an improvement of 30 seconds!

Even though these improvements might sound very attractive, sometimes changing the mirror mode is not an option. For instance, your company's disaster-recovery strategy can require a synchronous mirror configuration. Changing the mirror mode from synchronous to asynchronous should, in my opinion, be a last option (if it actually is a viable option). There are other parts that can influence DBMIRROR_SEND Wait Times, like the storage configuration on the mirror or the network connection between the principal and mirror SQL Server instances. Both these parts can act like a bottleneck between both instances, contributing to the DBMIRROR_SEND Wait Time.

Next to checking out the performance of your storage subsystem and network connection, SQL Server has a database mirroring monitor that will give you status information about the mirroring configuration. You can find the database mirroring monitor by right-clicking the database that is part of a mirror, selecting Tasks - database mirroring monitor. Figure 9-5 shows the monitor against my test mirror configuration.

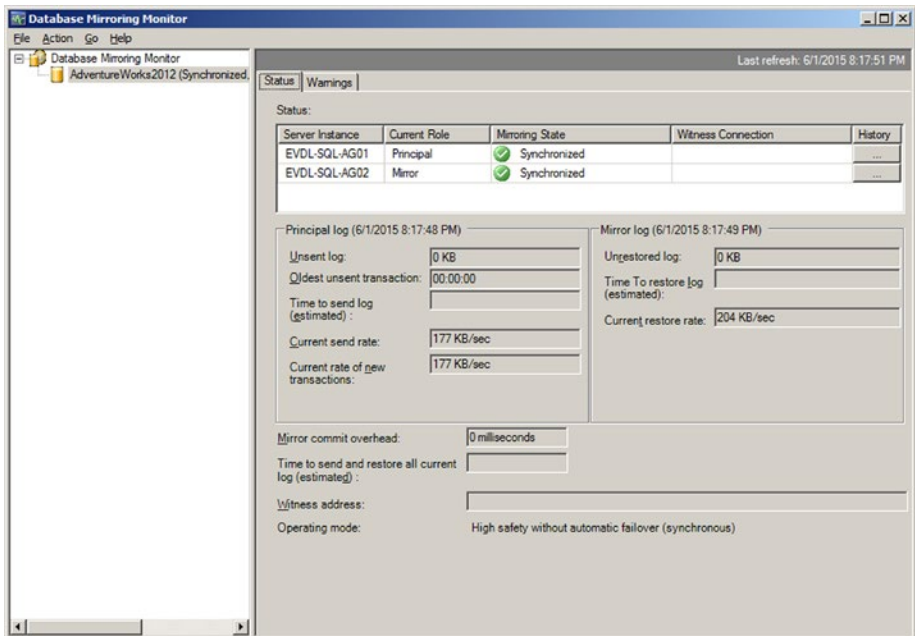


Figure 9-5. Database mirroring monitor

As you can see, the database mirroring monitor can provide you with some very interesting additional information like the number of log records that still need to be send or restored, how far behind the mirror currently is, and the send and restore rates. In many of my dealings with database mirroring, the database mirroring monitor is the first place I'll check when there are performance issues involving the mirror configuration.

DBMIRROR_SEND Summary

The DBMIRROR_SEND Wait Type is directly related to database mirroring. Seeing DBMIRROR_SEND waits occur is pretty normal on most mirror configurations. This makes using a baseline to identify Wait Times spikes a necessity. The mirroring mode plays a huge part in the DBMIRROR_SEND Wait Times. When using synchronous mirroring, DBMIRROR_SEND Wait Times will frequently be higher than when using asynchronous mirroring. Not only the mirroring mode influences DBMIRROR_SEND Wait Times, though. Having a storage subsystem on the mirror SQL Server instance that cannot keep up with the load will have an effect on DBMIRROR_SEND Wait Times, just like the network connection between the principal and the mirror.

HADR_LOGCAPTURE_WAIT & HADR_WORK_QUEUE

The HADR_LOGCAPTURE_WAIT and HADR_WORK_QUEUE Wait Types are both related to AlwaysOn Availability Groups. All Wait Types that are related to AlwaysOn can easily be identified by the HADR_ prefix in the Wait Type's name. AlwaysOn Availability Groups was introduced in SQL Server 2012 as a replacement for various SQL Server high-availability and disaster-recovery features such as database mirroring. There are quite a few different Wait Types associated with AlwaysOn, totaling 50 in SQL Server 2014. Not all of these Wait Types necessarily indicate performance problems somewhere in your AlwaysOn configuration. The HADR_LOGCAPTURE_WAIT and HADR_WORK_QUEUE are both perfect examples of benign Wait Types that occur naturally over time and do not directly indicate a performance problem. Since both these Wait Types have high Wait Times associated with them on every AlwaysOn configuration, and are thus very common, I wanted to include them inside this chapter to help you better understand what function they have.

What Are the HADR_LOGCAPTURE_WAIT & HADR_WORK_QUEUE Wait Types?

As I mentioned in the section above, both the HADR_LOGCAPTURE_WAIT and HADR_WORK_QUEUE Wait Types occur on AlwaysOn configurations. They both occur on different places inside your AlwaysOn configuration and have slightly different functions.

According to Books Online, the HADR_LOGCAPTURE_WAIT Wait Type indicates that SQL Server is “waiting for log records to become available. Can occur either when waiting for new log records to be generated by connections or for I/O completion when reading log not in the cache. This is an expected wait if the log scan is caught up to the end of log or is

reading from disk.” The `HADR_LOGCAPTURE_WAIT` Wait Type occurs on the SQL Server that hosts the primary database inside an AlwaysOn Availability Group. Think of the primary database as being just like the principal inside a database mirroring configuration.

AlwaysOn works much the same way as database mirroring and also provides two different modes (called Availability modes inside AlwaysOn): Synchronous-commit and Asynchronous-commit. Both these Availability modes work in the same way as their database mirroring counterparts we were discussing earlier in this chapter do. This means that in Synchronous-commit mode the primary replica waits to commit transactions to the transaction log until the secondary replica has completed its own log hardening, while in Asynchronous-commit mode the primary replica will directly commit the transaction to the transaction log without waiting for a confirmation from the secondary replica.

While the primary replica is waiting for work, SQL Server will record the time it has spent on waiting for new transactions to become available as the `HADR_LOGCAPTURE_WAIT` Wait Type. This means that seeing high Wait Times on the `HADR_LOGCAPTURE_WAIT` Wait Type actually means that SQL Server is waiting on new transactions to become available so they can be transferred to the secondary replica. This is not dependent on the Availability mode you configured for your AlwaysOn Availability Group. The `HADR_LOGCAPTURE_WAIT` Wait Type will always occur, no matter your AlwaysOn configuration. Figure 9-6 shows an AlwaysOn Availability Group configuration together with the `HADR_LOGCAPTURE_WAIT` Wait Type on the primary replica, which occurs while waiting for new transactions to be sent to the secondary replica.




Figure 9-6. AlwaysOn Availability Group and the `HADR_LOGCAPTURE_WAIT` Wait Type

Even though I placed the `HADR_LOGCAPTURE_WAIT` Wait Type on the primary replica in Figure 9-6, it will also log Wait Type on the secondary replica, although those values will normally be much lower than on the primary replica.

The `HADR_WORK_QUEUE` Wait Type is almost identical in function to the `HADR_LOGCAPTURE_WAIT` Wait Type. Books Online gives an excellent description of this Wait Type: “AlwaysOn Availability Groups’ background worker thread waiting for new work to be assigned. This is an expected wait when there are ready workers waiting for new work, which is the normal state.” The main difference between both Wait Types is

that the `HADR_LOGCAPTURE_WAIT` Wait Type is dedicated to waiting until new transactions become available, while the `HADR_WORK_QUEUE` indicates that are free threads waiting for work. Just like the `HADR_LOGCAPTURE_WAIT` Wait Type, the `HADR_WORK_QUEUE` occurs on both the primary and the secondary replicas, but the `HADR_WORK_QUEUE` Wait Type is much more prevalent on both replicas. As a matter of fact, the `HADR_WORK_QUEUE` Wait Type will frequently be the top AlwaysOn related Wait Type on every SQL Server that is part of an AlwaysOn Availability Group, especially if the work load is low.

Figure 9-7 shows an AlwaysOn Availability Group like the one in Figure 9-6, but this time I added the `HADR_WORK_QUEUE` Wait Type to the image as well.




Figure 9-7. `HADR_LOGCAPTURE_WAIT` and `HADR_WORK_QUEUE` Wait Types

Since both the `HADR_LOGCAPTURE_WAIT` and `HADR_WORK_QUEUE` Wait Types occur naturally over time, I did not include an example of both the Wait Types. Also, because both these Wait Types are not directly related to performance problems, there is no use including a section on lowering the Wait Times of both these Wait Types.

HADR_LOGCAPTURE_WAIT & HADR_WORK_QUEUE Summary

Both the `HADR_LOGCAPTURE_WAIT` and `HADR_WORK_QUEUE` Wait Types are benign Wait Types that occur on every SQL Server that is part of an AlwaysOn Availability Group. Because the `HADR_LOGCAPTURE_WAIT` and `HADR_WORK_QUEUE` Wait Types are not directly related to performance problems there is no direct need to focus attention on lowering them, and they can, in most cases, be safely ignored.

HADR_SYNC_COMMIT

The HADR_SYNC_COMMIT Wait Type is another AlwaysOn-related Wait Type that was introduced in SQL Server 2012. In many ways the HADR_SYNC_COMMIT Wait Type closely resembles the DBMIRROR_SEND Wait Type we discussed earlier in this chapter. There are some differences, however, between both Wait Types, which we will discuss in the section below.

What Is the HADR_SYNC_COMMIT Wait Type?

The HADR_SYNC_COMMIT Wait Type indicates the time the primary replica spends waiting for the secondary replica to harden the log records. HADR_SYNC_COMMIT waits will only occur on the primary replica and only inside a synchronous-replication AlwaysOn Availability Group. As soon as a transaction is received by the primary replica and is sent to the secondary replica for hardening, the HADR_SYNC_COMMIT Wait Time will start recording. The HADR_SYNC_COMMIT Wait Time will only stop recording when the secondary replica has sent its confirmation that the write to the secondary's transaction log was completed. Figure 9-8 below shows the HADR_SYNC_COMMIT Wait Time generation inside a timeline.




Figure 9-8. HADR_SYNC_COMMIT and synchronous replication

Since the HADR_SYNC_COMMIT Wait Type will always occur in every synchronous replicated AlwaysOn Availability Group, it is normal to expect a certain amount of Wait Time. But just like the DBMIRROR_SEND Wait Type, the Wait Time of the HADR_SYNC_COMMIT Wait Type is highly dependent on the speed at which the secondary replica can process the log records. This means that a slow network connection between both replicas or the performance of the storage subsystem on the secondary replica can impact the HADR_SYNC_COMMIT Wait Times. For this reason it is important to understand what the normal Wait Times for the HADR_SYNC_COMMIT Wait Type are for your AlwaysOn configuration so you can identify higher-than-normal Wait Times easier.

HADR_SYNC_COMMIT Example

For this example I have build an AlwaysOn Availability Group configured to use synchronous replication. The configuration of the test machines I will use for this can be found in Appendix I Example SQL Server Machine Configuration. I won't go into detail on how you can configure an AlwaysOn Availability Group, as there is plenty of information available on the Internet to help you configure AlwaysOn. A good starting point is the "Getting Started with AlwaysOn Availability Groups" article on Books Online, which you can find here: <https://msdn.microsoft.com/en-us/gg509118>. I used the AdventureWorks database as the database that needed to be replicated inside my AlwaysOn Availability Group.

After my AlwaysOn Availability Group was configured I added an extra table named AO_Test to the AdventureWorks database using the script in Listing 9-3.

Listing 9-3. Create AO_Test table

```
USE [AdventureWorks]
GO

CREATE TABLE AO_Test
(
    ID UNIQUEIDENTIFIER PRIMARY KEY,
    RandomData VARCHAR(50)
);
```

After the table is created I first clear and then query the sys.dm_os_wait_stats DMV to check the current Wait Times on the HADR_SYNC_COMMIT Wait Time on both the primary and secondary replica using the this query:

```
SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'HADR_SYNC_COMMIT';
```

Even after waiting for a couple of minutes the Wait Time of the HADR_SYNC_COMMIT Wait Type stays 0, as you can see in Figure 9-9. This is what I expected, since we have not performed any data modifications on the primary replica so far.

Primary

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	HADR_SYNC_COMMIT	0	0	0	0

Secondary

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	HADR_SYNC_COMMIT	0	0	0	0

Figure 9-9. HADR_SYNC_COMMIT wait information during no activity

This is a difference as compared to `HADR_LOGCAPTURE_WAIT` and `HADR_WORK_QUEUE`, which will accumulate Wait Times even though (or because) there is no user activity inside the AlwaysOn Availability Group.

Now that the table is in place, let's generate some transactions by performing a number of inserts. The script in Listing 9-4 will insert 10,000 rows into the `AO_Test` table we created earlier.

Listing 9-4. Insert 10,000 rows into the `AO_Test` table

```
INSERT INTO AO_Test
(
  ID,
  RandomData
)
VALUES
(
  NEWID(),
  CONVERT(VARCHAR(50), NEWID())
);
GO 10000
```

When the script in Listing 9-4 has completed I check the Wait Statistics information inside the `sys.dm_os_wait_stats` DMV again on both the primary and secondary replicas. Figure 9-10 shows the results of this query.

Primary

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	HADR_SYNC_COMMIT	10000	18439	57	232

Secondary

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	HADR_SYNC_COMMIT	0	0	0	0

Figure 9-10. `HADR_SYNC_COMMIT` waits on the primary and secondary replicas

The first thing you will notice when looking at Figure 9-10 is that the `HADR_SYNC_COMMIT` waits only occur on the primary replica and not on the secondary, which is expected behavior. The second interesting thing is the number of waits that occurred. This is the exact same amount as the number of rows we inserted. Again, this is expected behavior. Since we performed a single insert and just repeated it 10,000 times, every insert generated a single transaction-log record that needed to be replicated. Using the number of waits that occurred and the Wait Time, it is possible to calculate the average time it took for one insert operation to be committed on the replica. In this case it is 1.84 milliseconds (18439 / 10000) which is a pretty decent value.

Lowering HADR_SYNC_COMMIT Waits

Seeing HADR_SYNC_COMMIT waits occur does not necessarily mean there is a problem. HADR_SYNC_COMMIT waits will always occur whenever there are data modifications performed on your primary replica. They can indicate a problem if the Wait Times are much higher than you expect them to be when you compare them to your baseline measurements.

Changing the AlwaysOn operation mode to asynchronous replication will completely remove HADR_SYNC_COMMIT waits, but at the risk of losing data when a disaster occurs. Also, to reach your company’s disaster-recovery or high-availability needs, you frequently do not have the luxury of just changing the AlwaysOn operating mode, and I advise you not to change it just to lower HADR_SYNC_COMMIT Wait Times.

Thankfully, there are many different methods you can use to monitor the performance of your AlwaysOn Availability Group, including the AlwaysOn Dashboard, DMVs, and Perfmon counters.

You can open the AlwaysOn Dashboard by right-clicking your AlwaysOn Availability Group and selecting the “Show Dashboard” option. The AlwaysOn Dashboard, by default, gives you some general information, like the servers inside the Availability Group and the synchronization state, about your AlwaysOn Availability Group, as shown in Figure 9-11.

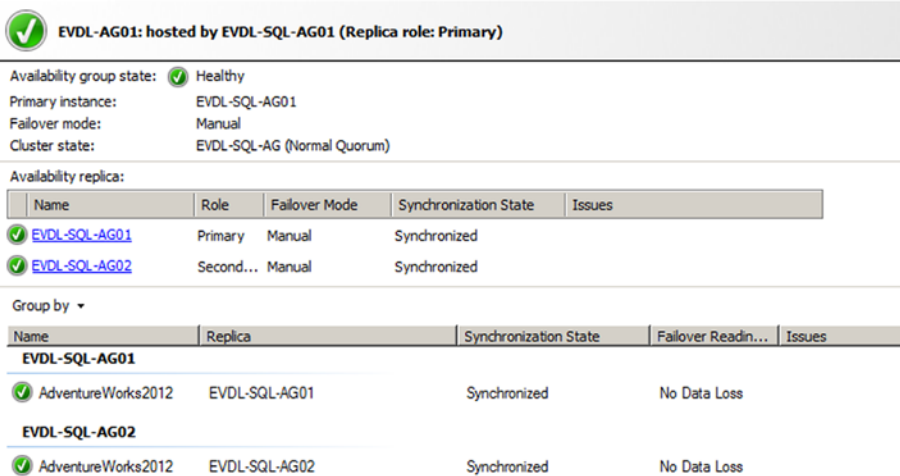


Figure 9-11. AlwaysOn Dashboard

The default view of the AlwaysOn Dashboard doesn’t provide much information you can use for troubleshooting. Thankfully you can configure the view to suit your own needs by right-clicking on the column bar and selecting the information you are interested in, as shown in Figure 9-12.

EVDL-AG01: hosted by EVDL-SQL-AG01 (Replica role: Primary)

Availability group state: ✔ Healthy
 Primary instance: EVDL-SQL-AG01
 Failover mode: Manual
 Cluster state: EVDL-SQL-AG (Normal Quorum)

Availability replica:

Name	Role	Failover Mode
✔ EVDL-SQL-AG01	Primary	Manual
✔ EVDL-SQL-AG02	Second...	Manual

Group by ▾

Name	Replica
EVDL-SQL-AG01	
✔ AdventureWorks2012	EVDL-SQL-AG01
EVDL-SQL-AG02	
✔ AdventureWorks2012	EVDL-SQL-AG02

Context menu items:

- Name
- Replica
- Synchronization State
- Failover Readiness
- Issues
- Suspended
- Suspend Reason
- Estimated Recovery Time (seconds)
- Estimated Data Loss (time)
- Synchronization Performance (seconds)
- Log Send Queue Size (KB)
- Log Send Rate (KB/sec)
- Redo Queue Size (KB)
- Redo Rate (KB/sec)

Figure 9-12. AlwaysOn add columns

There are many columns that are interesting for troubleshooting synchronization issues, and I recommend taking the time to understand them so you can determine which columns are most applicable to your situation.

The information shown by the AlwaysOn Dashboard is originally recorded inside various AlwaysOn-related DMVs. This makes it possible for you to query this information yourself. All of the AlwaysOn-related DMVs can easily be identified by the `dm_hadr` prefix in the DMV name, like the `sys.dm_hadr_database_replica_states` DMV that contains a large part of the information you can access inside the AlwaysOn Dashboard.

Next to the AlwaysOn and DMVs that are related to AlwaysOn, there are a large amount of Perfmon counters that specifically show AlwaysOn performance. These counters are grouped in the Perfmon `SQLServer:Availability Replica` and `SQLServer:Database Replica` groups. Figure 9-13 shows a part of the counters available in the `SQLServer:Database Replica` group.

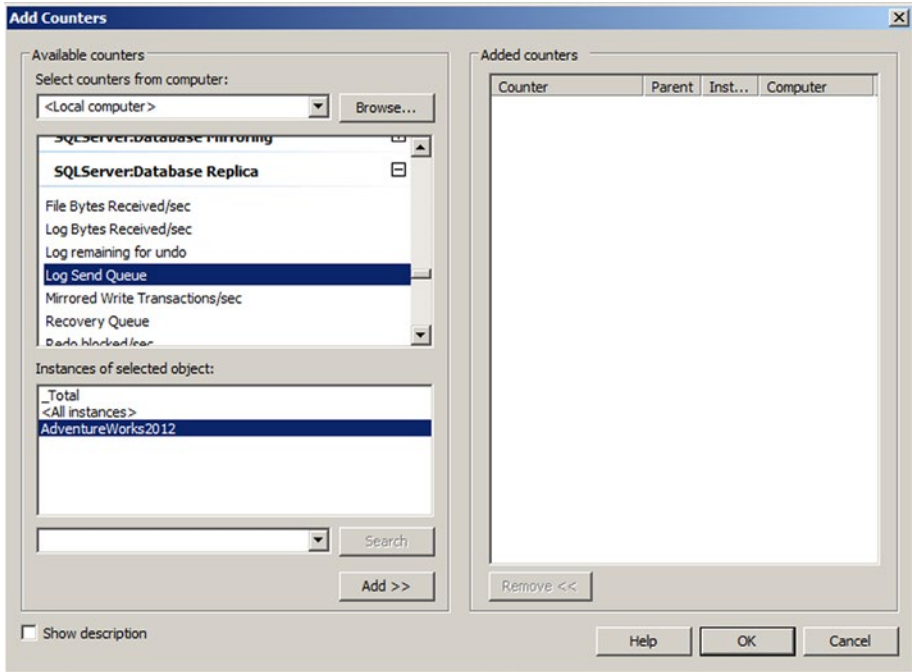


Figure 9-13. Perfmon counters related to AlwaysOn

As you have read so far, there are plenty of options available to you for analyzing the AlwaysOn performance between replicas.

Using the information from the various sources I have shown you so far, you should be able to check the general health of your AlwaysOn Availability Group. You can then combine this information with other metrics for things that impact the performance of your secondary replica, like the performance of your storage subsystem and your network connection. Since the HADR_SYNC_COMMIT Wait Type is strictly related to the secondary replica, you should focus your analysis on the SQL Server that hosts the secondary replica. For instance, if your storage subsystem cannot keep up with the amount of transactions that need to be committed on the secondary replica, you will notice this in higher HADR_SYNC_COMMIT Wait Times, and also in the various counters inside the AlwaysOn Dashboard, DMVs, or Perfmon.

It is difficult to give a general recommendation on how to lower HADR_SYNC_COMMIT Wait Times since they are highly dependent on myriad variables and also depend on your workload. When you have a workload that consists of a large amount of read queries you will notice lower HADR_SYNC_COMMIT Wait Times than workloads that perform many data modification operations. This means analyzing and optimizing your query workload can also contribute to the lowering of HADR_SYNC_COMMIT Wait Times.

HADR_SYNC_COMMIT Summary

The HADR_SYNC_COMMIT Wait Type will only occur on AlwaysOn Availability Groups that consist of replicas that are configured to use the synchronous replication mode. The HADR_SYNC_COMMIT Wait Type will give you insight into how long it took for the secondary replica to commit the transaction to disk. Since the HADR_SYNC_COMMIT will always record Wait Times inside synchronous replication, you should only worry about the Wait Times when they are far higher than expected. Thankfully there are various methods available to you to analyze the performance of your AlwaysOn Availability group, including an AlwaysOn Dashboard, DMVs, and Perfmon counters.

Since the performance of the secondary replica has the largest impact on the HADR_SYNC_COMMIT Wait Times, your attention should focus on the secondary replica when troubleshooting this Wait Type. The storage subsystem and network connection both play a large role in the speed at which the secondary replica can write log records to its transaction log. Your workload also impacts HADR_SYNC_COMMIT Wait Times, and optimizing it so data modifications are better spread out will result in lower HADR_SYNC_COMMIT Wait Times.

REDO_THREAD_PENDING_WORK

The last Wait Type in this chapter is the REDO_THREAD_PENDING_WORK Wait Type. And even though it misses the characteristic HADR_ prefix that identifies AlwaysOn-related Wait Types, it is related to AlwaysOn. The REDO_THREAD_PENDING_WORK Wait Type is, just like the HADR_LOGCAPTURE_WAIT and HADR_WORK_QUEUE Wait Types, a Wait Type that accumulates over time when there is no work to be done. And just like the HADR_LOGCAPTURE_WAIT and HADR_WORK_QUEUE Wait Types, it can in most cases be ignored since it does not indicate a performance problem.

Even though this is a Wait Type that can safely be ignored in 99% of the cases, I wanted to include it in this chapter for two reasons. It is usually one of the top Wait Types on an AlwaysOn Availability Group secondary replica, and understanding its related process inside SQL Server will give you a better understanding of the inner workings of AlwaysOn.

What Is the REDO_THREAD_PENDING_WORK Wait Type?

The REDO_THREAD_PENDING_WORK Wait Type is related to a process that only occurs on the secondary replica inside an AlwaysOn Availability Group, the Redo Thread.

Up to this point in the chapter, we have talked about how the secondary replica inside an AlwaysOn Availability Group processes log records, hardens them to its own transaction log, and sends a confirmation to the primary replica. When using synchronous replication, the primary replica will wait before sending a transaction complete message to the client that started the transaction, and when using asynchronous replication the message is sent without waiting for the hardening on the secondary. But until now we haven't discussed the process that will perform the modifications inside the secondary database described in the log records. This is where the Redo Thread on the secondary comes in. This thread is responsible for performing

the data modifications that were recorded in the log records the primary replica sent it. There is one very important concept associated with the Redo Thread: it does not impact the commit confirmation from the secondary replica. This means that the Redo Thread might be performing work long after the transaction has been communicated as committed to the client (both the primary and secondary replica have hardened the log record and the AlwaysOn Availability Group has the synchronized status).

This means that even though your AlwaysOn Availability Group is synchronized, the data inside the secondary database does not necessarily have to be identical to the primary database. This actually matters less than you might think on first thought. Because the secondary hardened the log records to its own transaction log on disk, it has all the information it needs to perform the Redo operation. Transactions will not be lost if a failure occurs on the primary, since the secondary has all the transactions that were performed in its own transaction log and can redo all the transactions. This works much the same as a standalone SQL Server instance where transactions are also hardened to disk first before data is actually changed. If SQL Server were to crash in this situation, SQL Server would use the transaction log to redo or undo the data modifications. Figure 9-14 shows an example of synchronous replication together with the Redo Thread. Note that the Redo Thread is a separate operation that does not impact the duration of the transaction complete message.

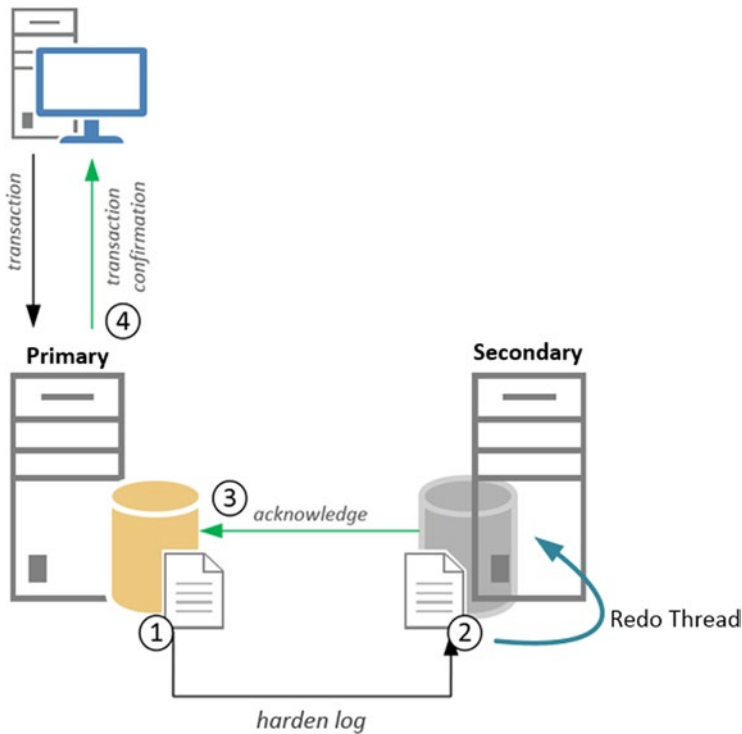


Figure 9-14. Synchronous AlwaysOn Availability Group and the Redo Thread

So, where does the REDO_THREAD_PENDING_WORK Wait Type come in? Well, if the Redo Thread is waiting for work to arrive, it will record the time it is inactive as Wait Time on the REDO_THREAD_PENDING_WORK Wait Type. This will occur on both synchronous and asynchronous replication modes, but only on the secondary replica.

Because the Wait Type only indicates that the Redo Thread is not performing any work, it can, save for extremely rare cases, be safely ignored. And because the Wait Time for the REDO_THREAD_PENDING_WORK Wait Type will accumulate naturally when there is no work to be done, there is no need to write an example demonstrating the Wait Type. A simple query to retrieve REDO_THREAD_PENDING_WORK Wait Type information against the `sys.dm_os_wait_stats` DMV on a secondary replica will show you that the Wait Time increases, especially when there is no user activity against the AlwaysOn Availability Group, as shown in Figure 9-15 below.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	REDO_THREAD_PENDING_WORK	170400	17168893	122	16006

Figure 9-15. REDO_THREAD_PENDING_WORK wait information

REDO_THREAD_PENDING_WORK Summary

The REDO_THREAD_PENDING_WORK Wait Type is an AlwaysOn-related Wait Type that accumulates Wait Time naturally over time when there is no data modification activity against an AlwaysOn Availability Group. The REDO_THREAD_PENDING_WORK Wait Type is related to the Redo Thread on the secondary replica inside an AlwaysOn Availability Group, and it indicates that the Redo Thread is currently waiting for work. Since this Wait Type will occur on every secondary replica, especially when there is minimal to no user data modification occurring, it can safely be ignored.

CHAPTER 10



Preemptive Wait Types

In Chapter 1, “Wait Statistics Internals,” we briefly touched upon SQL Server’s non-preemptive scheduling model that is used to perform thread scheduling and management. Unlike SQL Server, the Windows operating system uses preemptive scheduling to schedule and manage threads. Sometimes SQL Server has to use Windows functions to perform specific actions through the operating system, for instance, when checking Active Directory permissions. When this occurs, SQL Server will have to ask a thread from the Windows operating system, outside of SQL Server, thus making it impossible for SQL Server to manage that thread. While SQL Server is waiting for the preemptive thread inside the Windows operating system to complete, SQL Server will record a wait on a preemptive Wait Type. Figure 10-1 shows a graphical representation of this behavior.

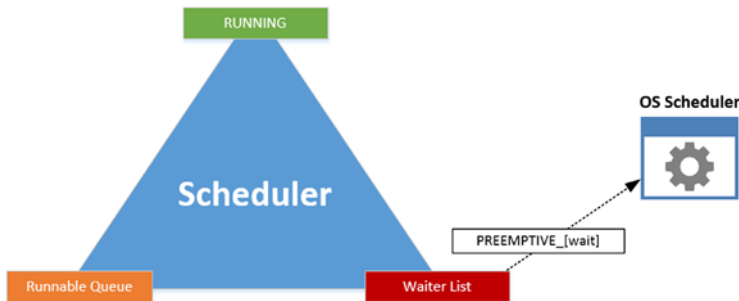


Figure 10-1. Preemptive wait occurring

There are many different preemptive Wait Types inside SQL Server; at the time of writing this book SQL Server 2014 has 192 different preemptive Wait Types. Which preemptive Wait Type is recorded when a thread is requested outside SQL Server depends on the Windows function the thread is accessing. Each of the preemptive Wait Types inside SQL Server represents a different Windows function (save for some exceptions that act as a catch-all Wait Type for different functions), and in many cases the name of the Wait Type is identical to the name of the Windows function. This is very helpful, because you can search for the specific Windows function on MSDN and

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-1-4842-1139-7_10](https://doi.org/10.1007/978-1-4842-1139-7_10)) contains supplementary material, which is available to authorized users.

learn what the function does. If you know what the function does, you also know why, or on what, SQL Server is waiting. For example, if you notice high Wait Times on the `PREEMPTIVE_OS_WRITEFILEGATHER` Wait Type you can remove the `PREEMPTIVE_OS_` part and search MSDN for the `WRITEFILEGATHER` function. Figure 10-2 shows the results I got on my search for the `WRITEFILEGATHER` function.

WriteFileGather function

Retrieves data from an array of buffers and writes the data to a file.

The function starts writing data to the file at a position that is specified by an `OVERLAPPED` structure. The `WriteFileGather` function operates asynchronously.

Figure 10-2. *WriteFileGather Windows function*

By reading the article we can learn a lot about this function; apparently this function is used when writing data to a file and has to occur outside SQL Server. I won't spoil anything else here, since we will go into more detail about the `PREEMPTIVE_OS_WRITEFILEGATHER` Wait Type a bit further down in this chapter.

I won't describe every possible preemptive Wait Type in this chapter, since there are simply too many of them. Instead I have focused on the most common preemptive Wait Types. If you run into a preemptive Wait Type that is not discussed in detail in this chapter, I suggest you use the method above to find more information about the Windows function on MSDN. Hopefully, that information can help you to figure out why the wait is occurring.

PREEMPTIVE_OS_ENCRYPTMESSAGE & PREEMPTIVE_OS_DECRYPTMESSAGE

The first preemptive Wait Types we are going to discuss in this chapter are the `PREEMPTIVE_OS_ENCRYPTMESSAGE` and `PREEMPTIVE_OS_DECRYPTMESSAGE` Wait Types. As you can probably guess from the Wait Type names, the functions are related to either encrypting or decrypting messages through the Windows operating system.

What Are the `PREEMPTIVE_OS_ENCRYPTMESSAGE` & `PREEMPTIVE_OS_DECRYPTMESSAGE` Wait Types?

As I noted in the previous section, the `PREEMPTIVE_OS_ENCRYPTMESSAGE` and `PREEMPTIVE_OS_DECRYPTMESSAGE` Wait Types are related to the encryption and decryption of messages. More specifically, they are related to encrypting and decrypting network traffic to and from the SQL Server instance. One case where this is used is when connecting to your SQL Server instance using certificates to encrypt the data that is sent between the client and the SQL Server instance. In that case, SQL Server will need to access the Windows operating system to perform the encryption of the messages that

it is sending to the client, or to decrypt the messages that are received. The encryption and decryption do not happen inside SQL Server, unlike, for instance, Transparent Data Encryption (TDE), where the encryption/decryption process happens entirely inside SQL Server.

Both the `PREEMPTIVE_OS_ENCRYPTMESSAGE` and `PREEMPTIVE_OS_DECRYPTMESSAGE` Wait Types do not necessarily indicate any performance problems. They just show you encryption is being used, so there is no real need to troubleshoot these Wait Types. The overhead of encrypting and decrypting messages is so small that it rarely causes any serious issues (I have yet to come across a case where using certificates to connect to SQL Server caused performance problems).

PREEMPTIVE_OS_ENCRYPTMESSAGE & PREEMPTIVE_OS_DECRYPTMESSAGE Example

To show you an example of both the `PREEMPTIVE_OS_ENCRYPTMESSAGE` and `PREEMPTIVE_OS_DECRYPTMESSAGE` Wait Types, I am going to configure a certificate that will be used to encrypt the connection to the SQL Server instance. To make this example reproducible, I have included the steps for creating a self-signed certificate. Normally, in production environments you will use a certificate issued by a certificate authority, but for testing purposes a self-signed certificate is fine.

The first thing I did to be able to generate a self-signed certificate is install Internet Information Services (IIS) on my test virtual machine. IIS makes generating a self-signed certificate very simple.

After the installation of IIS is completed, I open the IIS Manager from Administrative Tools. Then I click on the name of my machine and select the Server Certificates option in the Features View, as shown in Figure 10-3 below.




Figure 10-3. Features View inside the IIS Manager

This will open a new Server Certificates view inside the IIS Manager. Inside the Action Pane, I click the Create Self-Signed Certificate option. I am then asked to supply a name for my certificate, so I filled in the name of my test virtual machine as you can see in Figure 10-4.

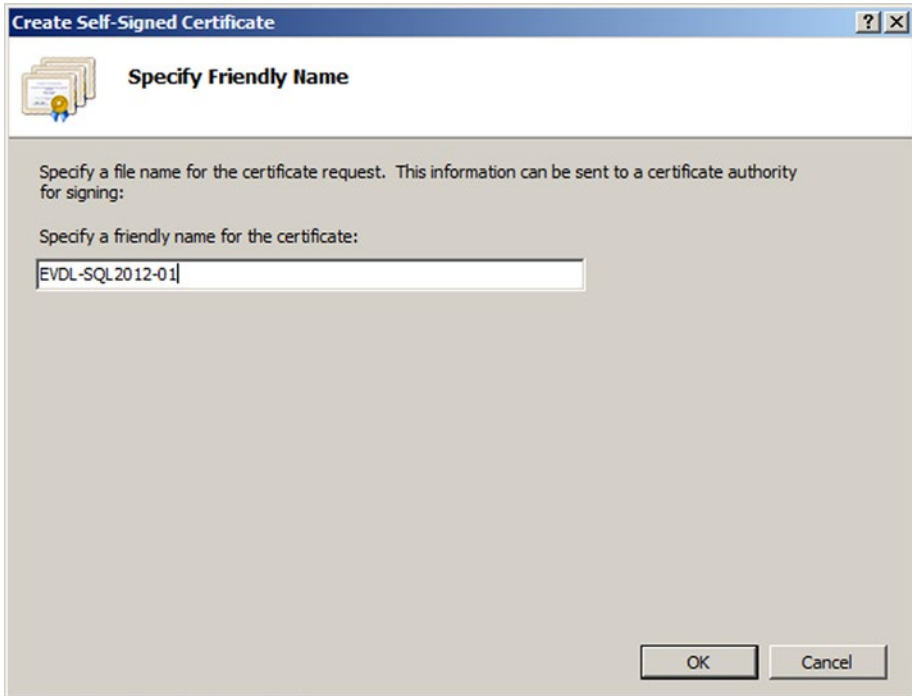


Figure 10-4. Create self-signed certificate

I click OK, and the self-signed certificate will be created and automatically placed inside the correct certificate store on my machine (Local Machine - Personal Certificates).

Now that my self-signed certificate is created and stored inside the certificate store, I need to make sure the account my SQL Server service is running under has permissions to access the certificates. I open MMC by clicking Start - Run, entering MMC, and pressing OK. Now that the MMC console is open I need to add the Certificates snap-in. I do this by clicking File - Add/Remove Snap-in, selecting the Certificates snap-in, and clicking Add. When prompted for which account I want to manage certificates, I select Computer account, as shown in Figure 10-5, and click Next and Finish.




Figure 10-5. Certificate account selection

Inside the Certificates console, I open the folder Certificates (Local Computer) - Personal - Certificates. If the generation of the self-signed certificate inside IIS was correct, I should see the certificate here. Figure 10-6 shows the certificate on my test virtual machine.

Issued To ^	Issued By	Expiration Date	Intended Purposes	Friendly Name
EVDL-SQL2012-01	EVDL-SQL2012-01	6/5/2016	Server Authentication	EVDL-SQL2012-01

Figure 10-6. self-signed certificate

I right-click the self-signed certificate and select All Tasks - Manage Private Keys. A permissions dialog opens. Here I need to add the account under which the SQL Server Service is running. In my case I use the local network account. The account only needs read permission on the certificate, as shown in Figure 10-7.

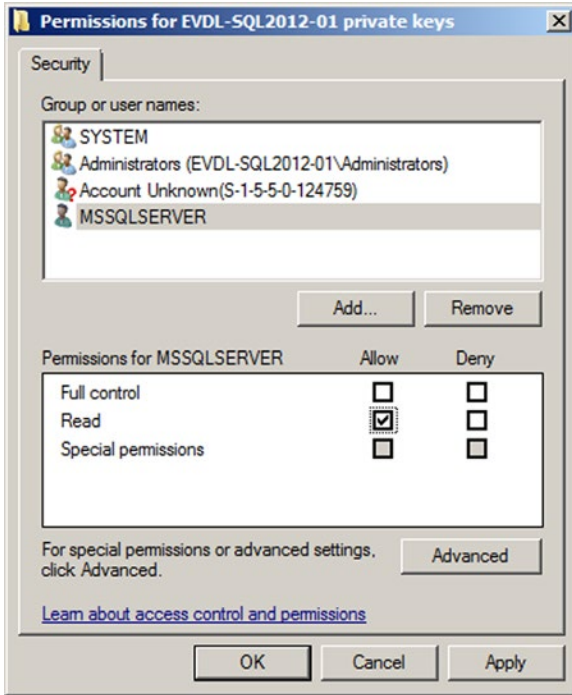


Figure 10-7. *Self-signed certificate permissions*

After adding the account and selecting the right permission, I click OK to close the dialog. Now that the permissions are correct and the SQL Server Service account can access the certificate, I need to add the self-signed certificate to the network configuration of the SQL Server instance that I want to enable for encryption. I open up the SQL Server Configuration Manager and click the SQL Server Network Configuration option. I right-click on the SQL Server instance that should use the self-signed certificate and select Properties. Next, I open the Certificates tab and select the self-signed certificate we created earlier. Figure 10-8 shows the dialog on my test virtual machine.

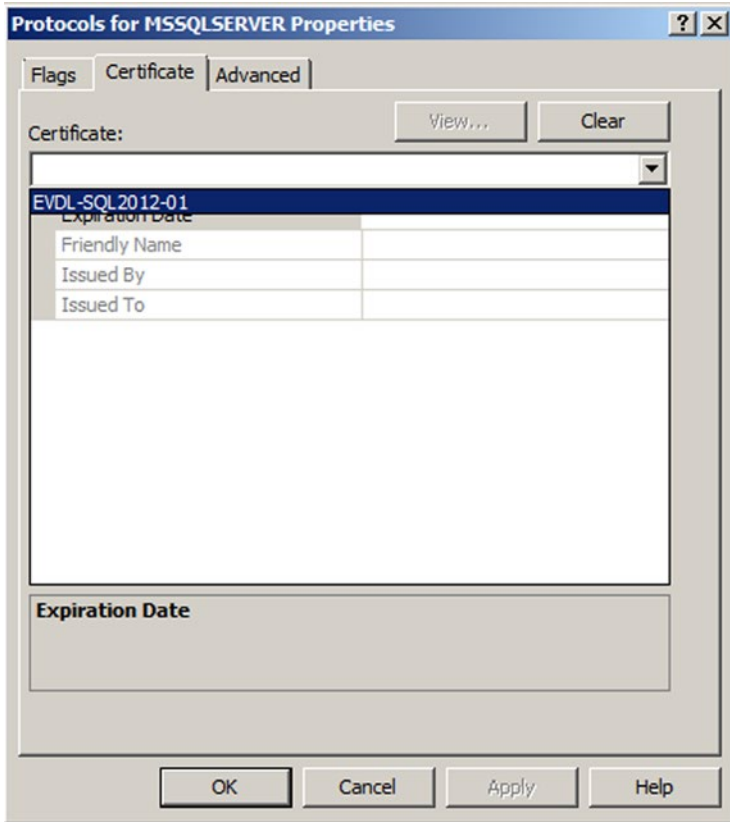


Figure 10-8. Certificate selection

After selecting the self-signed certificate, I click OK to close the dialog. I am notified that the certificate will become active after a restart of the SQL Server service, so I perform a restart of the SQL Server Service.

Right now SQL Server can use the self-signed certificate, but to make sure my network messages are encrypted I have to connect to the SQL Server instance and tell it that I want to use encryption. For this example I will use the SQL Server Management Studio, on the same virtual machine as my SQL Server instance, to connect to the SQL Server instance. If you connect to your SQL Server instance from another machine you need to make sure the self-signed certificate is available on that machine. When the Connect to Server dialog appears inside SQL Server Management Studio I click the Options button at the bottom right of the dialog. This opens up additional properties for the connection to my SQL Server instance. I select the Encrypt connection checkbox as shown in Figure 10-9, and connect to my SQL Server instance.



Figure 10-9. Connection Properties in SQL Server Management Studio

Right now I have configured everything I need to make sure SQL Server will use the self-signed certificate to encrypt messages between the SQL Server instance and SQL Server Management Studio, so I can finally take a look at the `PREEMPTIVE_OS_ENCRYPTMESSAGE` and `PREEMPTIVE_OS_DECRYPTMESSAGE` Wait Types!

Generating the `PREEMPTIVE_OS_ENCRYPTMESSAGE` and `PREEMPTIVE_OS_DECRYPTMESSAGE` waits is very simple now. Basically, every query I execute from the SQL Server Management Studio right now will be encrypted, even if I run SQL Server Management Studio on the same machine as the SQL Server instance. I use the query in Listing 10-1 to reset the `sys.dm_os_wait_stats` DMV, connect to the AdventureWorks database, perform a simple query, and then look at the waits occurring on the `PREEMPTIVE_OS_ENCRYPTMESSAGE` and `PREEMPTIVE_OS_DECRYPTMESSAGE` Wait Types.

Listing 10-1. Select query using encrypted connection

```
DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR)

USE AdventureWorks
GO

SELECT *
FROM Sales.SalesOrderDetail;

SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'PREEMPTIVE_OS_ENCRYPTMESSAGE'
OR wait_type = 'PREEMPTIVE_OS_DECRYPTMESSAGE';
```

The results of these queries on my test SQL Server instance are shown in Figure 10-10.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	PREEMPTIVE_OS_DECRYPTMESSAGE	1	0	0	0
2	PREEMPTIVE_OS_ENCRYPTMESSAGE	2381	339	3	0

Figure 10-10. *PREEMPTIVE_OS_DECRYPTMESSAGE & PREEMPTIVE_OS_ENCRYPTMESSAGE waits*

As you can see, the `PREEMPTIVE_OS_ENCRYPTMESSAGE` Wait Time has more waits and Wait Time associated with it. This is logical since I performed a select query and it only had to decrypt the acknowledgment network messages from the client. The results of the query had to be encrypted by SQL Server, which leads to higher waits on the `PREEMPTIVE_OS_ENCRYPTMESSAGE` Wait Type.

Lowering `PREEMPTIVE_OS_ENCRYPTMESSAGE` & `PREEMPTIVE_OS_DECRYPTMESSAGE` Waits

Under normal circumstances there should be no need to focus attention on lowering the `PREEMPTIVE_OS_ENCRYPTMESSAGE` and `PREEMPTIVE_OS_DECRYPTMESSAGE` Wait Types. They mostly just indicate that message encryption is occurring, which is probably a choice that was made when configuring the SQL Server instance. Disabling encryption will dramatically lower the `PREEMPTIVE_OS_ENCRYPTMESSAGE` and `PREEMPTIVE_OS_DECRYPTMESSAGE` Wait Times, but at the cost of security.

`PREEMPTIVE_OS_ENCRYPTMESSAGE` & `PREEMPTIVE_OS_DECRYPTMESSAGE` Summary

The `PREEMPTIVE_OS_ENCRYPTMESSAGE` and `PREEMPTIVE_OS_DECRYPTMESSAGE` Wait Types indicate that encryption is occurring between the SQL Server instance and a client. These Wait Types can generally be ignored since they do not directly indicate a performance problem. Lowering them can be achieved by disabling the use of encryption, but this comes at the cost of security.

PREEMPTIVE_OS_WRITEFILEGATHER

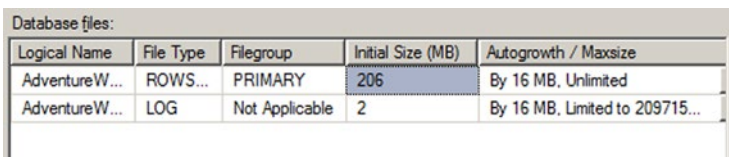
The `PREEMPTIVE_OS_WRITEFILEGATHER` Wait Type is related to storage interactions, more specifically the writing to files through the Windows operating system.

What Is the `PREEMPTIVE_OS_WRITEFILEGATHER` Wait Type?

The `PREEMPTIVE_OS_WRITEFILEGATHER` Wait Type is related to the `WriteFileGather` function inside the Windows operating system. If we look up the definition of this function on Books Online we get the following description: “Retrieves data from an array of buffers and writes the data to a file.” From this description we can assume the function will be called when there is a need to write data to a file. This does not count for every storage subsystem write operation inside SQL Server, however. Generally there is no need for SQL Server to move outside of its own engine to wait for a preemptive operation. There are some exceptions, however, which can result in `PREEMPTIVE_OS_WRITEFILEGATHER` waits (depending on the Windows function used to perform the storage subsystem interaction). One specific operation inside SQL Server that will always result in `PREEMPTIVE_OS_WRITEFILEGATHER` waits is the growing of data files. Whenever SQL Server wants to grow a data file it will need to allocate extra space on the storage subsystem and “zero out” the new space so SQL Server can use it. The allocation of the extra space does not happen inside the SQL Server Engine, thus a preemptive operation has to take place, which can lead to preemptive waits on the `WriteFileGather` function.

`PREEMPTIVE_OS_WRITEFILEGATHER` Example

To show you an example of `PREEMPTIVE_OS_WRITEFILEGATHER` waits occurring, I am going to replicate the situation I described in the previous section, growing a database data file. For this example I will restore a backup of the AdventureWorks database; in the SQL Server 2014 version of the database there exists only a single database data file with a size of 206 MB, as shown in Figure 10-11.



Database files:				
Logical Name	File Type	Filegroup	Initial Size (MB)	Autogrowth / Maxsize
AdventureW...	ROWS...	PRIMARY	206	By 16 MB, Unlimited
AdventureW...	LOG	Not Applicable	2	By 16 MB, Limited to 209715...

Figure 10-11. Default database file configuration of AdventureWorks (2014 edition)

I am going to grow the single database data file to a size of 10 GB. Because the allocation of the extra space needed for the data file is performed outside SQL Server, this should result in `PREEMPTIVE_OS_WRITEFILEGATHER` waits.

To perform the action of enlarging the database data file I used the script shown in Listing 10-2. This script will clear the `sys.dm_os_wait_stats` DMV, enlarge the AdventureWorks data file to 10 GB, and then query the `sys.dm_os_wait_stats` DMV for `PREEMPTIVE_OS_WRITEFILEGATHER`.

Listing 10-2. Enlarge AdventureWorks database data file

```
DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR)

USE [master]
GO

ALTER DATABASE [AdventureWorks]
    MODIFY FILE
    (
        NAME = N'AdventureWorks_Data',
        SIZE = 10485760KB
    );
GO

SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'PREEMPTIVE_OS_WRITEFILEGATHER';
```

The query in Listing 10-2 takes around 21 seconds on my test SQL Server instance and results in the wait information shown in Figure 10-12 for the `PREEMPTIVE_OS_WRITEFILEGATHER` Wait Type.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	PREEMPTIVE_OS_WRITEFILEGATHER	1	21465	21465	0

Figure 10-12. `PREEMPTIVE_OS_WRITEFILEGATHER` waits

Notice that there was only one single wait on the `PREEMPTIVE_OS_WRITEFILEGATHER` Wait Type, and the duration was practically as long as it took to perform the enlargement of the data file.

Lowering `PREEMPTIVE_OS_WRITEFILEGATHER` Waits

When you notice higher-than-normal Wait Times on the `PREEMPTIVE_OS_WRITEFILEGATHER` Wait Type it means that a process from inside SQL Server is performing actions on the storage subsystem through the Windows operating system. The first matter of action should be to investigate what process initiated the action that resulted in `PREEMPTIVE_OS_WRITEFILEGATHER` waits. Very frequently this will be the (automatic) growth of a database data or log file. If you allow the data or log files to grow automatically when they are full, you can expect to see `PREEMPTIVE_OS_WRITEFILEGATHER` waits occur whenever an auto-growth

event occurs. This does not necessarily mean there is a problem, but if auto-growth events take a long time to complete because, for instance, the storage subsystem is experiencing performance problems, your queries might experience performance degradation as well.

There is one Windows setting available that I frequently see not configured, instant file initialization. We discussed this setting, and how you can enable it, already in Chapter 5, “IO-Related Wait Types,” under the `ASYNC_IO_COMPLETION` Wait Type so I won’t go into detail on how to enable the setting again. Figure 10-13 shows the results of the query in Listing 10-2 with instant file initialization enabled, and as you can see, the amount of Wait Time spent on the `PREEMPTIVE_OS_WRITEFILEGATHER` Wait Time disappeared completely. The amount of time it took to execute the query from Listing 10-2 also went down to less than one second, instead of the 21 seconds without instant file initialization.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	PREEMPTIVE_OS_WRITEFILEGATHER	0	0	0	0

Figure 10-13. *PREEMPTIVE_OS_WRITEFILEGATHER* waits with instant file initialization turned on

Next to using instant file initialization, the performance of the storage subsystem plays a large part in the `PREEMPTIVE_OS_WRITEFILEGATHER` Wait Times. The better your storage subsystem performs, the lower the Wait Times on the `PREEMPTIVE_OS_WRITEFILEGATHER` Wait Type.

Another SQL Server action that can cause higher-than-normal `PREEMPTIVE_OS_WRITEFILEGATHER` Wait Times is performing database restores. Much like expending a data file, before SQL Server can restore a database it needs to allocate free storage for it. This is also related to instant file initialization, which will also speed up database restores just like file enlargements.

PREEMPTIVE_OS_WRITEFILEGATHER Summary

The `PREEMPTIVE_OS_WRITEFILEGATHER` Wait Type indicates that SQL Server is asking the Windows operating system to perform an operation on the storage subsystem. Not all operations can be handled from inside the SQL Server Engine and actions; for example, the growing of a data file requires the execution of a Windows function to allocate the desired space on the storage subsystem. Instant file initialization is a setting in Windows that can lower the amount of `PREEMPTIVE_OS_WRITEFILEGATHER` Wait Time drastically, but the performance of the storage subsystem itself also plays a large role in `PREEMPTIVE_OS_WRITEFILEGATHER` Wait Times.

PREEMPTIVE_OS_AUTHENTICATIONOPS

The `PREEMPTIVE_OS_AUTHENTICATIONOPS` Wait Type is another preemptive Wait Type that is related to various Windows authentication functions.

What Is the `PREEMPTIVE_OS_AUTHENTICATIONOPS` Wait Type?

The `PREEMPTIVE_OS_AUTHENTICATIONOPS` Wait Type is recorded whenever SQL Server needs to perform an account authentication, for instance, to authenticate the SQL Server Windows login when it connects to SQL Server. Seeing `PREEMPTIVE_OS_AUTHENTICATIONOPS` waits occur is to be expected, especially when using mixed-mode authentication and Windows logins inside your SQL Server instance.

One common misconception about the `PREEMPTIVE_OS_AUTHENTICATIONOPS` Wait Type is that it is only related to SQL Server logins that use Windows authentication inside a domain. This is not entirely correct. While it is true that `PREEMPTIVE_OS_AUTHENTICATIONOPS` Wait Times will frequently be higher when using Active Directory accounts to connect to SQL Server, `PREEMPTIVE_OS_AUTHENTICATIONOPS` waits will also occur if the SQL Server instance is installed on a machine outside of a domain; the Wait Times will generally be lower though.

Figure 10-14 shows a simplified image of how SQL Server connects to an Active Directory domain controller to validate the SQL Server Windows login. Keep in mind that the Windows operating system takes care of the communication between the domain controller and the SQL Server, hence the preemptive Wait Type.




Figure 10-14. SQL Server Windows login authentication inside domain

On a machine that has a SQL Server instance installed but is not part of a domain, the authentication of the Windows login will occur on the machine itself (local accounts).

Because it will generally take a longer time to authenticate a Windows login through a domain controller (the request has to travel across the network and authenticate on another machine), the Wait Times for the `PREEMPTIVE_OS_AUTHENTICATIONOPS` Wait Type will generally be higher for a SQL Server instance inside a domain.

PREEMPTIVE_OS_AUTHENTICATIONOPS Example

To generate an example of `PREEMPTIVE_OS_AUTHENTICATIONOPS` waits occurring, I do not need to perform any complex actions. Opening a new connection to the SQL Server instance using Windows authentication should be enough. One way to make this easy to measure is by connecting to the SQL Server instance with SQL Server Management Studio, using Windows authentication. Figure 10-15 shows the SQL Server Management Studio connect dialog to my test SQL Server instance. Note that my test SQL Server instance is not inside a domain, and that I use the local administrator account on my machine to connect.



Figure 10-15. connect SQL Server Management Studio using local Windows authentication

The next step I perform is opening a new Query Window inside SQL Server Management Studio and performing the steps inside the query shown in Listing 10-3.

Listing 10-3. Generate `PREEMPTIVE_OS_AUTHENTICATIONOPS` waits

```
-- Step 1 Clear sys.dm_os_wait_stats
DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR);

-- Step 2 Open a new Query Window inside
-- SQL Server Management Studio

-- Step 3 go back to this Query Window
-- and run the query below
```



```
SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'PREEMPTIVE_OS_AUTHENTICATIONOPS';
```

If you follow the steps commented inside the script in Listing 10-3 you should see `PREEMPTIVE_OS_AUTHENTICATIONOPS` waits occurring after running the query in step 3. Figure 10-16 shows the results of the query in step 3 on my test machine.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	PREEMPTIVE_OS_AUTHENTICATIONOPS	12	2	1	0

Figure 10-16. *PREEMPTIVE_OS_AUTHENTICATIONOPS* waits

As you can see, the number of waits occurring and their Wait Times are very low. The point of this example is not to show you an example of very high `PREEMPTIVE_OS_AUTHENTICATIONOPS` Wait Times, but rather how they occur naturally when connecting to a SQL Server instance. Because I opened a new Query Window inside SQL Server Management Studio, a new connection to the SQL Server instance will be made using the Windows login I used to connect to my SQL Server instance. Because it is a new connection, the account I used to connect had to be authenticated, resulting in `PREEMPTIVE_OS_AUTHENTICATIONOPS` waits.

Lowering `PREEMPTIVE_OS_AUTHENTICATIONOPS` Waits

In the example, I have shown you how `PREEMPTIVE_OS_AUTHENTICATIONOPS` waits occur naturally whenever you connect to a SQL Server instance. Now imagine a situation where your SQL Server instance is part of a domain environment and it uses Windows authentication to authenticate domain users (or groups) against an Active Directory. In that case your authentication request has to travel across the network in order to perform the authentication of the account. There are many factors involved that can impact the speed of the authentication request; for instance, if your Domain Controller is under a lot of stress it can take longer to perform the authentication, or if your network experiences performance degradation it will also impact the authentication request. These factors also contribute to the `PREEMPTIVE_OS_AUTHENTICATIONOPS` Wait Type and can result in higher Wait Times.

I would like to describe to you a case I encountered at a client that involved the `PREEMPTIVE_OS_AUTHENTICATIONOPS` Wait Type to give you an idea of how you can lower `PREEMPTIVE_OS_AUTHENTICATIONOPS` Wait Times.

At this client they used an application that connected to SQL Server using the Windows account that was logged in on the computer that ran the application. The computers and the SQL Server were all part of a domain. From a security perspective the application was well designed, as it did not require separate SQL Server users who needed permission on the database and also didn't use a generic account to connect to the SQL Server instance and execute queries. Inside, the database-specific objects (like tables) were also secured based on domain users and groups.

The client started to experience server performance problems inside the application after deploying it to every (3000+) computer inside the company. The DBA at the client couldn't find any problems, there were no infrastructure-related performance problems on the SQL Server instance, and executing the queries on the SQL Server instance itself revealed no issues. When we looked at the Wait Statistics we noticed that the most prevalent Wait Type was the `PREEMPTIVE_OS_AUTHENTICATIONOPS` Wait Type. We also noticed that the application would connect to the SQL Server instance, run a query, then disconnect again. Because so many concurrent users were using the application, it resulted in a high amount of Windows authentication requests, so many that the Domain Controller couldn't handle them, resulting in the slower processing of authentication requests.

In this case the Domain Controller was a virtual machine, and after adding more processor and memory resources it was able to keep up with the high amount of authentication requests.

As you can see from this case, seeing high `PREEMPTIVE_OS_AUTHENTICATIONOPS` Wait Times does not necessarily mean your SQL Server instance is running into problems, especially in a domain environment, as the performance of your Domain Controller also plays a large role in `PREEMPTIVE_OS_AUTHENTICATIONOPS` Wait Times.

The moral of the story is, if you notice higher-than-normal `PREEMPTIVE_OS_AUTHENTICATIONOPS` Wait Times you will need to investigate much more than the SQL Server instance. Make sure to check the performance of your Domain Controllers if you are using Windows authentication inside a domain. Check every infrastructure part between your SQL Server instance and the Domain Controller, like network switches, firewalls, and so on. All of these infrastructure parts will add additional latency for each authentication request, which will result in higher `PREEMPTIVE_OS_AUTHENTICATIONOPS` Wait Times, making it a difficult Wait Type to troubleshoot.

PREEMPTIVE_OS_AUTHENTICATIONOPS Summary

The `PREEMPTIVE_OS_AUTHENTICATIONOPS` Wait Type is related to performing authentication requests by the Windows operating system. It is normal to see `PREEMPTIVE_OS_AUTHENTICATIONOPS` waits occur, especially when your SQL Server instance is part of a domain and uses Windows authentication to authenticate users. Higher-than-normal Wait Times can indicate that authentication requests are taking longer than normal to complete. This does not necessarily mean that your SQL Server instance is running into a performance problem. If the Domain Controller cannot process the authentication requests fast enough, it will result in higher `PREEMPTIVE_OS_AUTHENTICATIONOPS` Wait Times. A slow network connection to the Domain Controller, firewall, or switch configurations can also impact `PREEMPTIVE_OS_AUTHENTICATIONOPS` Wait Times.

PREEMPTIVE_OS_GETPROCADDRESS

The final Wait Type I would like to discuss in this chapter is the `PREEMPTIVE_OS_GETPROCADDRESS` Wait Type. The `PREEMPTIVE_OS_GETPROCADDRESS` Wait Type is related to the execution of extended stored procedures inside SQL Server.

Extended stored procedures allow you to create external routines in a language other than T-SQL; for instance, using the C# programming language. These extended stored procedures are loaded into SQL Server using .dll files and can expand the capabilities of SQL Server programming by allowing you to perform actions that would be impossible in T-SQL, like reading/writing Windows Registry entries.

Extended stored procedures are marked as deprecated since SQL Server 2008 and Common Language Runtime (CLR) should be used instead of them. However, there are still cases inside SQL Server that require extended stored procedures, and some third-party software vendors still rely on them.

What Is the `PREEMPTIVE_OS_GETPROCADDRESS` Wait Type?

The `PREEMPTIVE_OS_GETPROCADDRESS` Wait Type is recorded whenever the entrypoint inside an extended stored procedure is loaded. The entrypoint is called upon whenever SQL Server loads or unloads the extended stored procedure .dll file. Under normal conditions the loading of the entrypoint should complete quickly, resulting in very low `PREEMPTIVE_OS_GETPROCADDRESS` Wait Times, if any, but depending on the extended stored procedure, or problems related to loading the extended stored procedure .dll file, it is possible to notice higher Wait Times. Important to keep in mind is that the `PREEMPTIVE_OS_GETPROCADDRESS` Wait Type only records the time it took to load the entrypoint of the .dll file, not the execution time of the extended stored procedure. Figure 10-17 shows a (simplified) overview of how extended stored procedures are executed by SQL Server.

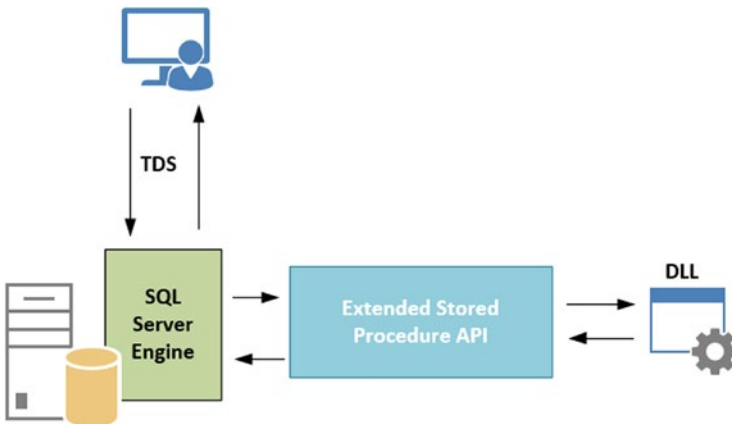


Figure 10-17. Executing an extended stored procedure

Not only can you write your own extended stored procedures to perform actions not possible using T-SQL, SQL Server itself comes shipped with many different extended stored procedures. Most of these can be recognized by the `xp_` prefix in the extended stored procedure name. Figure 10-18 shows a selection of extended stored procedures inside the master database of my test SQL Server instance.

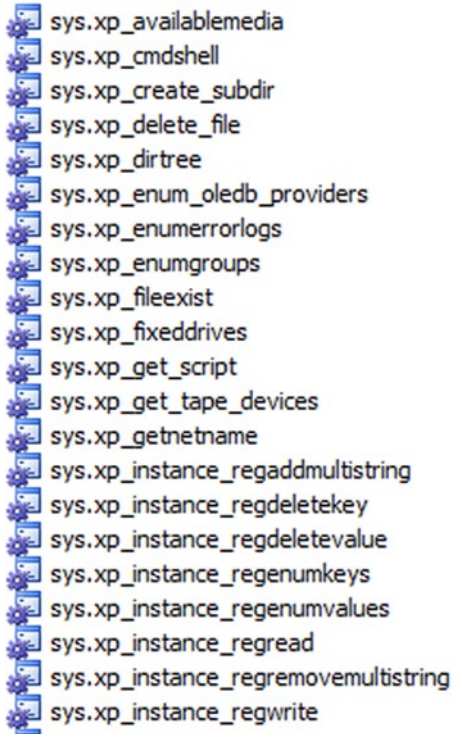


Figure 10-18. Selection of extended stored procedures inside the master database

Probably the most notorious extended stored procedure is the `xp_cmdshell` extended stored procedure. The `xp_cmdshell` extended stored procedure makes it possible to execute a command inside a Windows command shell from within SQL Server. This is a huge security risk if your SQL Server instance is compromised, since it gives access to commands that can affect the entire Windows operating system. Thankfully, it is impossible to run the `xp_cmdshell` extended stored procedure by default; you have to specifically allow its use by configuring an advanced configuration setting.

PREEMPTIVE_OS_GETPROCADDRESS Example

For this example I will execute an extended stored procedure already present inside SQL Server, `xp_getnetname`, instead of writing a custom extended stored procedure, which is far beyond the scope of this book. The `xp_getnetname` is an undocumented extended stored procedure that returns the NETBIOS name of the machine that hosts your SQL Server instance. Before executing the `xp_getnetname` extended stored procedure I clear the `sys.dm_os_wait_stats` DMV, and after the execution of `xp_getnetname` I query the DMV for `PREEMPTIVE_OS_GETPROCADDRESS` wait information. Listing 10-4 shows the entire query I executed on my test SQL Server instance.

Listing 10-4. Execute `xp_getnetname` and query Wait Statistics

```
USE [master]
GO

DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR);

exec xp_getnetname;

SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'PREEMPTIVE_OS_GETPROCADDRESS';
```

The results of the query in Listing 10-4 can be seen in Figure 10-19.

Server Net Name	
1	EVDL-SQL2012-01

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	PREEMPTIVE_OS_GETPROCADDRESS	1	0	0	0

Figure 10-19. `PREEMPTIVE_OS_GETPROCADDRESS` wait

The results aren't spectacular. Apparently, the `xp_getnetname` extended stored procedure doesn't cause any problems when loading the `.dll` entrypoint, since there is no Wait Time recorded. A wait still did occur though, as you can see in the `waiting_tasks_count` column, it just took SQL Server less than a millisecond to load the entrypoint.

Lowering `PREEMPTIVE_OS_GETPROCADDRESS` Waits

Since the `PREEMPTIVE_OS_GETPROCADDRESS` Wait Type is directly related to executing extended stored procedures, the first step in your investigation should be to detect what extended stored procedure is being executed and what its function is.

I have seen `PREEMPTIVE_OS_GETPROCADDRESS` waits occur at a number of clients because they were using a third-party backup application that used extended stored procedures to perform a database backup, but there are many more possible causes for high `PREEMPTIVE_OS_GETPROCADDRESS` Wait Times. Knowing which extended stored

procedure is being executed can help you trace what process is executing the extended stored procedure.

There have also been some known bugs inside SQL Server 2008 and 2008R2 that reported higher-than-normal `PREEMPTIVE_OS_GETPROCADDRESS` Wait Times because the execution time of the extended stored procedure was also recorded in the Wait Times, instead of only the entrypoint loading. If you are still using SQL Server 2008 or 2008R2 and experience very high `PREEMPTIVE_OS_GETPROCADDRESS` Wait Times, it might be worth your while to upgrade to the latest Service Pack and check if the `PREEMPTIVE_OS_GETPROCADDRESS` Wait Times go down.

PREEMPTIVE_OS_GETPROCADDRESS Summary

The `PREEMPTIVE_OS_GETPROCADDRESS` Wait Type is directly related to the execution of extended stored procedures. Extended stored procedures can be written in a variety of programming languages like C# and allow you to perform actions that would otherwise be impossible in T-SQL. Wait Time for the `PREEMPTIVE_OS_GETPROCADDRESS` Wait Type is recorded whenever the entrypoint inside an extended stored procedure .dll is loaded. In normal situations Wait Times for the `PREEMPTIVE_OS_GETPROCADDRESS` Wait Type are very low. Seeing high `PREEMPTIVE_OS_GETPROCADDRESS` Wait Times can indicate that the entrypoint loading is running into problems. There have also been bugs related to the calculation of the `PREEMPTIVE_OS_GETPROCADDRESS` Wait Type inside SQL Server 2008 and 2008R2. If you are running SQL Server 2008 or 2008R2 and experience high `PREEMPTIVE_OS_GETPROCADDRESS` Wait Times it might be worth your time to upgrade to the latest Service Pack.



Background and Miscellaneous Wait Types

SQL Server has many different internal processes that can run into a wait that results in a specific Wait Type, and so far we have discussed quite a few of them. Some of these internal processes are constantly running inside SQL Server, waiting until there is work for them to do. While these processes, frequently called *background processes*, are waiting for work to arrive, SQL Server will record the time they are waiting for work as Wait Time on specific Wait Types related to these background processes. While these background Wait Types are not directly related to performance problems, they frequently have the highest Wait Time and will show up at the top of the Wait Time list when you query the top Wait Types ordered by Wait Time.

Frequently these background Wait Types are called *benign*, and can safely be ignored because they simply indicate that an internal process is waiting for work to arrive. This logic is also true for the Wait Types we will discuss in this chapter, but instead of just telling you to ignore them when analyzing Wait Statistics, I want to give you some background information about them so you know what they measure and why it is safe to ignore them. Keep in mind that we are still talking about SQL Server here, and that means that “it depends” on many factors as to whether you can completely ignore these background Wait Types. You wouldn’t be the first person to run into a performance problem and only to find out that an ignored background process was actually the cause of the issue. So my advice is to ignore but to not forget about them!

Next to the background Wait Types I also added a number of miscellaneous Wait Types that were difficult to place in an earlier chapter because they didn’t quite fit in with the chapter’s Wait Type category.

Since the background Wait Types inside this chapter record Wait Time constantly when their associated processes are waiting for work to do, I did not include an example section or a lowering Wait Time section for these Wait Types.

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-1-4842-1139-7_11](https://doi.org/10.1007/978-1-4842-1139-7_11)) contains supplementary material, which is available to authorized users.

CHECKPOINT_QUEUE

The first Wait Type in this chapter is one of those background Wait Types that accumulates large amounts of Wait Time over time CHECKPOINT_QUEUE. The CHECKPOINT_QUEUE Wait Type can in many cases be safely ignored, but understanding what the Wait Type stands for and why it has such high Wait Times can't hurt.

What Is the CHECKPOINT_QUEUE Wait Type?

The CHECKPOINT_QUEUE Wait Type is related to the checkpoint process in SQL Server that is responsible for writing “dirty” (modified) data pages from the buffer cache to the data file on disk. In Chapter 5, “IO-Related Wait Types,” we took a good look at the checkpoint process when we discussed the SLEEP_BPOOL_FLUSH Wait Type, so I won't repeat all the information again here. What is important to know, and the reason why this Wait Type can normally be ignored, is that the CHECKPOINT_QUEUE Wait Type indicates that the checkpoint process is waiting for work. This means that Wait Times on the CHECKPOINT_QUEUE Wait Type don't indicate any performance issues; they just indicate the time the checkpoint processes spent waiting on work. On SQL Server instances that aren't very busy, or don't see many data modification operations, the Wait Time can reach very high values.

The recording of CHECKPOINT_QUEUE Wait Times inside the `sys.dm_os_wait_stats` and `sys.dm_os_waiting_tasks` DMVs goes through a specific internal routine that might return unexpected Wait Times (like sudden spikes inside your baseline). Figure 11-1 shows the results of queries against the `sys.dm_os_wait_stats` and `sys.dm_os_waiting_tasks` DMVs for wait information of the CHECKPOINT_QUEUE Wait Type on my test SQL Server instance.

wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1 CHECKPOINT_QUEUE	0	0	0	0

waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address
1 0x000000017F01D0C8	15	0	105413896	CHECKPOINT_QUEUE	0x000000003794EF70

Figure 11-1. CHECKPOINT_QUEUE waits

What is interesting to notice here is that the cumulative Wait Times inside the `sys.dm_os_wait_stats` DMV stay at 0, while the Wait Times inside the `sys.dm_os_waiting_tasks` DMV have a very high value. My test SQL Server instance doesn't perform much work in the background, so it is logical that the checkpoint process spends most of its time waiting for work. The reason for the difference in Wait Times between both DMVs is related to the way SQL Server executes checkpoint operations. The Wait Times shown in both the DMVs are only recorded by the automatic checkpoint process. A manual checkpoint execution does not impact the Wait Times. As part of the automatic checkpoint process, the Wait Times of the `sys.dm_os_waiting_tasks` DMV are moved to the `sys.dm_os_wait_stats` DMV and reset to 0. So, if you notice very high CHECKPOINT_QUEUE Wait Times inside the `sys.dm_os_waiting_tasks`, it means it was some time ago that the automatic checkpoint process ran.

To show you a simple demonstration of this behavior, I created a table, reset the `sys.dm_os_wait_stats` DMV, inserted a few rows inside the table, performed a manual checkpoint, and queried the `sys.dm_os_wait_stats` and `sys.dm_os_waiting_tasks` DMVs, as shown in Listing 11-1.

Listing 11-1. CHECKPOINT_QUEUE example

```
-- Create a table in the AdventureWorks database
USE [AdventureWorks]
GO

CREATE TABLE check_test
(
  ID UNIQUEIDENTIFIER,
  RandomData VARCHAR(50)
);
GO

-- Clear sys.dm_os_wait_stats
DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR);

-- Insert a few rows into our table
INSERT INTO check_test
(
  ID,
  RandomData
)
VALUES
(
  NEWID(),
  CONVERT(varchar(50), NEWID())
);
GO 100

CHECKPOINT 1;

-- Query Wait Statistics
SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'CHECKPOINT_QUEUE';

SELECT *
FROM sys.dm_os_waiting_tasks
WHERE wait_type = 'CHECKPOINT_QUEUE';
```

Figure 11-2 shows the results of the queries made against the `sys.dm_os_wait_stats` and `sys.dm_os_waiting_tasks` DMVs.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	CHECKPOINT_QUEUE	0	0	0	0

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address
1	0x000000017F01D0C8	15	0	107514985	CHECKPOINT_QUEUE	0x000000003794EF70

Figure 11-2. CHECKPOINT_QUEUE waits

As you can see, the manual checkpoint didn't generate any waits inside the `sys.dm_os_wait_stats` DMV. Also, an automatic checkpoint didn't occur, because inserting 100 rows generated too few log records to trigger an automatic checkpoint.

If we were to insert more rows, we should be able to trigger an automatic checkpoint. In this case I ran the query below to insert 100,000 rows into the table we created in Listing 11-1. While the insert was running, I queried the `sys.dm_os_wait_stats` and `sys.dm_os_waiting_tasks` DMVs repeatedly to see if anything changed. See the following:

```
INSERT INTO check_test
(
  ID,
  RandomData
)
VALUES
(
  NEWID(),
  CONVERT(varchar(50), NEWID())
);
GO 100000
```

After a few seconds I noticed that the Wait Time for the CHECKPOINT_QUEUE Wait Type was moved to the `sys.dm_os_wait_stats` DMV, as shown in Figure 11-3.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	CHECKPOINT_QUEUE	1	107881628	107881628	0

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address
1	0x000000017F01D0C8	15	0	384	CHECKPOINT_QUEUE	0x000000003794EF70

Figure 11-3. CHECKPOINT_QUEUE waits

Apparently we inserted enough log records to cause an automatic checkpoint to occur, and as you can see from this example, only an automatic checkpoint will write the Wait Times of the CHECKPOINT_QUEUE to the `sys.dm_os_wait_stats` DMV.

You should keep this behavior in mind when you notice sudden, very high Wait Time values inside the `sys.dm_os_wait_stats` DMV. This will normally only occur in SQL Server instances that either have a very small workload or that have a workload mainly consisting of read operations instead of data modification operations.

CHECKPOINT_QUEUE Summary

The CHECKPOINT_QUEUE Wait Type is related to checkpoint operations inside SQL Server. Wait Time on the CHECKPOINT_QUEUE Wait Type is recorded while SQL Server is waiting for an automatic checkpoint operation to take place. This is one of the Wait Types you can normally safely ignore because it doesn't indicate there are any performance issues. The Wait Times on the CHECKPOINT_QUEUE Wait Type are recorded differently between the `sys.dm_os_wait_stats` and `sys.dm_os_waiting_tasks` DMV, and this can cause sudden high Wait Times when querying the `sys.dm_os_wait_stats` DMV. Keep this behavior in mind when noticing high CHECKPOINT_QUEUE Wait Times inside the `sys.dm_os_wait_stats` DMV.

DIRTY_PAGE_POLL

The DIRTY_PAGE_POLL Wait Type was introduced in SQL Server 2012 with the indirect checkpoint feature and behaves a lot like the previous Wait Type we discussed, CHECKPOINT_QUEUE. While the automatic checkpoint process runs at a set interval of one minute, the indirect checkpoint feature allows you to configure a specific checkpoint interval on a per-database basis. Even if you are not using indirect Checkpoint, the DIRTY_PAGE_POLL Wait Type will still accumulate Wait Time.

What Is the DIRTY_PAGE_POLL Wait Type?

The DIRTY_PAGE_POLL Wait Type is another background wait that can normally be safely ignored. The Wait Type is related to the recovery writer process that is used by the indirect checkpoint feature that runs continuously in the background of your SQL Server instance. Because of this connection, let's take a quick look what indirect Checkpoints are and how they work.

As we know, the checkpoint process inside SQL Server is responsible for writing modified data pages from the buffer cache to the database data file on disk. By default, the checkpoint process runs automatically every minute, or when enough log records have been generated. The checkpoint process plays a vital part in the recovery duration of your SQL Server databases when a crash occurs. Take, for instance, the following scenario: While you are performing many modifications to a database in your SQL Server instance, a crash occurs. Luckily, you were able to simply restart the SQL Server service to get everything up and running again. The first thing SQL Server will do is start a recovery process. The recovery process will check the transaction log for any transactions that were not committed when the crash occurred and perform a rollback of the transaction. The recovery process will also check whether any data pages that were modified by a committed transaction received their modification inside the database data file because they were only modified in the buffer cache.

If any of those pages are found, SQL Server will use the transaction log to redo these transactions. Now imagine you have a busy SQL Server instance where many thousands of modifications are performed every minute. This means that the chance that there is a high number of dirty pages not written to disk yet is pretty high. If your SQL Server then crashes (or, for instance, a failover occurs), the recovery process will take more time to complete. Indirect checkpoints can help us keep this recovery process as short as

possible. By configuring this feature we can tell SQL Server to write modified data pages to disk faster; for instance, every 10 seconds. Figure 11-4 shows the location and name of the indirect checkpoint feature inside the properties of a database.

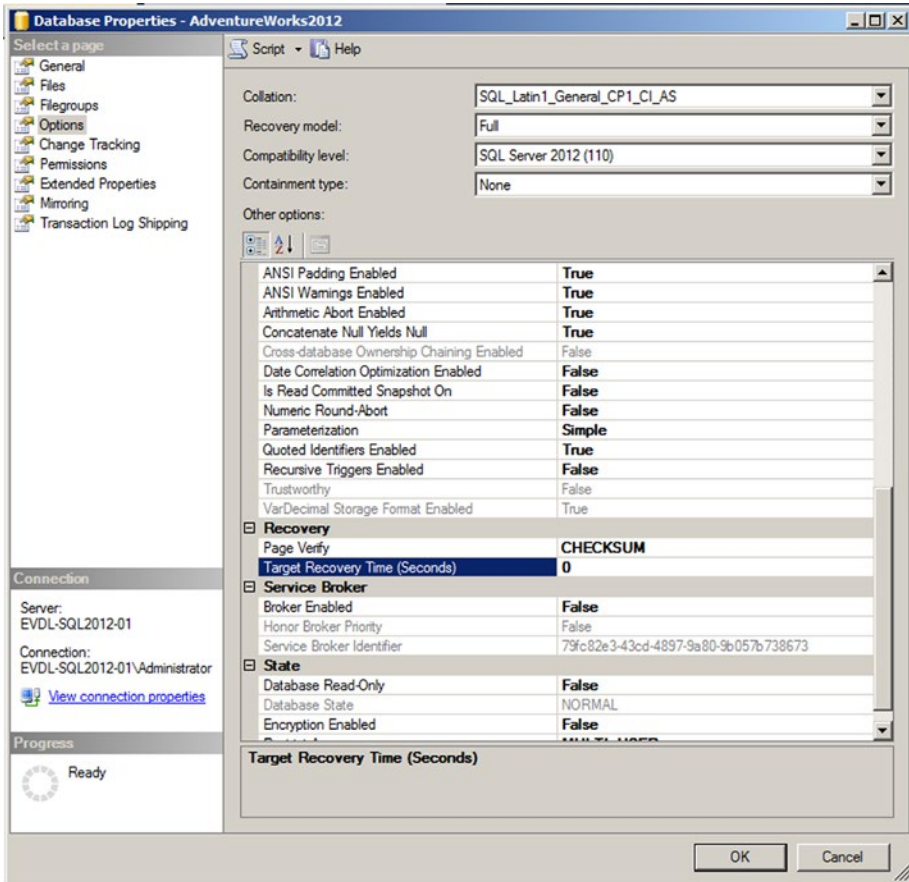


Figure 11-4. Indirect checkpoint feature location and value

By default, the value of the **Target Recovery Time (Seconds)** configuration option is 0. This means that indirect checkpoints are not being used. If you modify the value to anything other than 0, an indirect checkpoint will occur at the interval in seconds you specified.

The time you configure in the **Target Recovery Time (Seconds)** option does not mean that every x seconds the checkpoint process will be executed, however. By setting this value SQL Server will calculate how many dirty pages can exist before they need to be written to the database data file so that the recovery process never takes longer than the time specified. So, for instance, if you configure the **Target Recovery Time (Seconds)** option to 15 seconds, SQL Server will write dirty pages to the database data file at such an interval that when the SQL Server instance fails it can be recovered within 15 seconds.

To monitor how many dirty pages are inside the buffer cache so SQL Server knows when the dirty-page threshold has been reached, the recovery writer was introduced. Even if you do not configure the **Target Recovery Time (Seconds)** option, DIRTY_PAGE_POLL waits will still occur because the recovery writer process will still poll the number of dirty pages inside the buffer cache, even though no action is taken upon that number. As you can see in Figure 11-5, the Wait Times can reach high values easily even when not using indirect checkpoints.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	DIRTY_PAGE_POLL	73090	7361796	128	3747

Figure 11-5. DIRTY_PAGE_POLL waits

Indirect checkpoints also have a risk associated with them. Configuring the **Target Recovery Time (Seconds)** option to a very low value can lead to extra load on the storage subsystem because dirty pages are continuously written to disk. Be sure to test the setting extensively before configuring it on your production SQL Server instances.

DIRTY_PAGE_POLL Summary

The DIRTY_PAGE_POLL Wait Type was introduced in SQL Server 2012 with the introduction of the indirect checkpoints feature. Even if you do not use indirect checkpoints, the DIRTY_PAGE_POLL Wait Type will still accumulate Wait Time because of the new recovery writer process. Normally the DIRTY_PAGE_POLL Wait Type does not indicate a performance problem, and as such it can safely be ignored when analyzing Wait Statistics on your SQL Server instance.

LAZYWRITER_SLEEP

The LAZYWRITER_SLEEP Wait Type is, surprise, related to the SQL Server internal lazywriter process. The lazywriter process shares some similarities with the checkpoint process we discussed earlier in this chapter, in that it also writes dirty pages from the buffer cache to the database data file. The similarities end here, though, because the reason why the lazywriter process writes these pages to the database data file is completely different than the checkpoint process.

What Is the LAZYWRITER_SLEEP Wait Type?

Just like with other Wait Types we have discussed so far in this chapter, the LAZYWRITER_SLEEP Wait Type occurs when an internal SQL Server process, in this case the lazywriter process, is waiting for work. The lazywriter process is a background process that will become active at a certain time interval. When it becomes active it will scan the size of the buffer cache and determine if there are enough free pages inside the buffer cache.

It is important that there are always a certain number of free pages inside the buffer cache so that new page requests can fit directly without first having to swap out other pages. If the lazywriter process determines there are enough free pages in the buffer cache, it will go back to sleep again and record the LAZYWRITER_SLEEP Wait Type while it is sleeping. However, if there are not enough free pages inside the buffer cache, the lazywriter process will detect, between checkpoints, which dirty pages in the buffer cache haven't been accessed for a while, write them to the database data file, and remove them from the buffer cache. So, if there are more than enough free pages inside the buffer cache, the lazywriter process doesn't have much work to do. If your SQL Server instance is under memory pressure, the lazywriter process will be far busier while swapping out dirty pages and freeing up room inside the buffer cache. Figure 11-6 will show the relationship of the checkpoint and lazywriter processes with a flowchart.

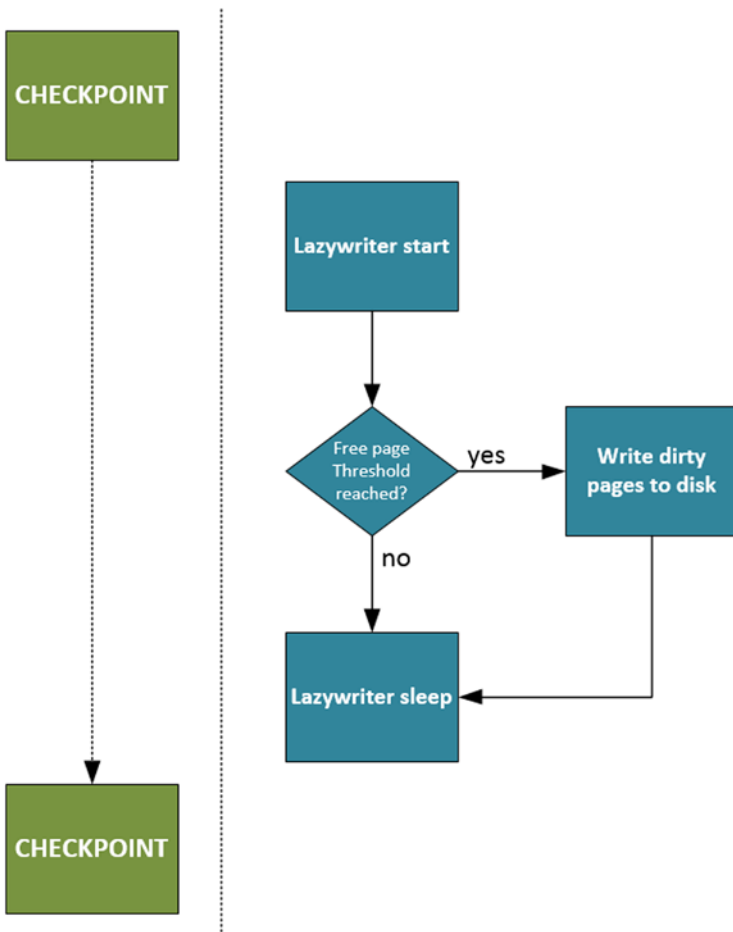


Figure 11-6. Checkpoint and lazywriter processes

Because the `LAZYWRITER_SLEEP` Wait Type indicates the time the lazywriter process spends sleeping, or waiting for work, it is another one of those Wait Types you can safely ignore. There is a catch, however—if the lazywriter process is constantly working to move dirty pages from the buffer cache to the database data file, it can indicate your SQL Server instance is experiencing memory pressure. This is bad for performance because every page has to be moved to the buffer cache before it can get read or modified. This behavior can potentially result in lower-than-normal Wait Times on the `LAZYWRITER_SLEEP` Wait Type.

LAZYWRITER_SLEEP Summary

The `LAZYWRITER_SLEEP` Wait Type is related to the lazywriter internal SQL Server process. The lazywriter process starts at a fixed time interval and is responsible for writing dirty data pages to the database data file if there are not enough free pages available inside the buffer cache. The `LAZYWRITER_SLEEP` Wait Type indicates that the lazywriter process is currently not running, or is sleeping, until it is signaled to wake up and check the buffer cache. Because the `LAZYWRITER_SLEEP` Wait Type only shows us how much time the lazywriter process spends being inactive, it can in most cases be ignored.

MSQL_XP

In the last section of the previous chapter, we discussed the `PREEMPTIVE_GETPROCADDRESS` Wait Type. We learned that the `PREEMPTIVE_GETPROCADDRESS` Wait Type records Wait Time when the entrypoint of an extended stored procedure is loaded. One important thing I noted was that the `PREEMPTIVE_GETPROCADDRESS` Wait Type does not record the execution time of the extended stored procedure, only the entrypoint loading. The execution time of an extended stored procedure is actually tracked by another Wait Type, `MSQL_XP`.

What Is the MSQL_XP Wait Type?

The `MSQL_XP` Wait Type records the execution time of extended stored procedures on your SQL Server instance. The `MSQL_XP` Wait Type is also used to detect deadlock situations when using Multiple Active Result Sets (MARS). MARS is a feature that allows the execution of multiple (concurrent) batches through a single SQL Server connection. We won't go further into detail about MARS, but you can find some more information about it here: <https://msdn.microsoft.com/en-us/library/ms131686.aspx>.

The most common reason for seeing higher-than-normal Wait Times on the `MSQL_XP` Wait Type is the execution of extended stored procedures. This does not necessarily mean there is a problem as long as the execution time of the extended stored procedures stays the same. However, if an extended stored procedure takes more time than expected, you are sure to notice it in the increase of the `MSQL_XP` Wait Time when comparing the Wait Time against a baseline.

MSQL_XP Example

To demonstrate that MSQL_XP waits occur when extended stored procedures are being executed, I created a simple example using the script in Listing 11-2. The script will reset the `sys.dm_os_wait_stats` DMV, execute an extended stored procedure (in this case the `xp_dirtree` extended stored procedure inside the master database), and query the `sys.dm_os_wait_stats` for MSQL_XP wait information.

Listing 11-2. MSQL_XP example

```
DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR);

EXEC master..xp_dirtree 'c:\windows';

SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'MSQL_XP';
```

The results of the query in Listing 11-2 on my test SQL Server instance can be seen in Figure 11-7. The top window shows the `xp_dirtree` results, the bottom window the results of the query against the `sys.dm_os_waits_stats` DMV.

	subdirectory	depth
1	AppCompat	1
2	Programs	2
3	AppPatch	1
4	AppPatch64	2
5	Custom	2
6	Custom64	3
7	en-US	2
8	assembly	1
9	GAC	2
10	ADODB	3
11	70.3300.0_h03f5f711d50a3a	4

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	MSQL_XP	1	6394	6394	0

Figure 11-7. MSQL_XP wait

The information in the `sys.dm_os_wait_stats` DMV shows us that one wait occurred on the MSQL_XP Wait Type with a Wait Time of 6394 milliseconds. This is almost identical to the time it took to execute the query in Listing 11-2, which was 6 seconds on my test SQL Server instance.

Lowering MSQL_XP Waits

When noticing higher-than-normal Wait Times for the MSQL_XP Wait Type, chances are that extended stored procedures are being used and are taking longer than normal to complete. Your first point of action should be to identify which extended stored procedures are being used and what they are being used for. Because extended stored procedures can also perform tasks outside SQL Server, they can run into other Windows processes that can slow them down. Knowing what the extended stored procedure function is, and what it does, can help you quickly identify where it is running into issues.

If you are using MARS, you are probably running into MARS-connection deadlocks. There have been various SQL Server updates that reduce the chances of MARS deadlocks occurring, so make sure your SQL Server instance is patched. Also make sure to check the application code that executes queries using MARS for potential issues.

MSQL_XP Summary

The MSQL_XP Wait Type does two different things: it detects the time it takes to execute extended stored procedures, and serves as deadlock detection for MARS connections. Seeing higher-than-normal Wait Times on the MSQL_XP Wait Type frequently indicates an extended stored procedure is taking longer than normal to complete. Try to detect which extended stored procedure is being executed and what its function is, as this will make troubleshooting the extended stored procedure easier.

OLEDB

The OLEDB Wait Type occurs whenever SQL Server has to access the Object Linking and Embedding Database (OLEDB) Client Provider. There are various reasons why SQL Server will use the OLEDB Client Provider, and whenever it does SQL Server will record Wait Time on the OLEDB Wait Type.

What Is the OLEDB Wait Type?

SQL Server uses the OLEDB Client Provider for many different actions inside SQL Server. For instance, linked server traffic will move through the OLEDB Client Provider and will result in OLEDB waits. Other actions, especially when SQL Server has to retrieve data from an outside source, can also result in OLEDB Client Provider usage.

Some actions inside SQL Server will also use the OLEDB Client Provider, even though they occur internally. One good example of this is the DBCC command, which I will demonstrate in the example section below.

OLEDB Example

One interesting process that uses the OLEDB Client Provider is the DBCC command inside SQL Server. Whenever you execute a DBCC command you are bound to see OLEDB waits occur. Listing 11-3 shows an example of OLEDB waits occurring after a DBCC CHECKDB. The example script will clear the `sys.dm_os_wait_stats` DMV, perform a CHECKDB against the AdventureWorks database, and then query the `sys.dm_os_wait_stats` DMV for OLEDB waits.

Listing 11-3. Generate OLEDB waits

```
DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR);

DBCC CHECKDB('AdventureWorks');

SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'OLEDB';
```

The results of the query in Listing 11-3 as performed against my test SQL Server instance can be seen in Figure 11-8.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	OLEDB	684132	334827	496	0

Figure 11-8. OLEDB waits

As you can see from Figure 11-8, performing a DBCC CHECKDB will lead to OLEDB waits.

Lowering OLEDB Waits

As you could see in the previous example, performing a DBCC CHECKDB against a database will result in OLEDB waits. This doesn't mean there is a problem related to the OLEDB Client Provider, however; rather, it just indicates that the DBCC CHECKDB command makes use of the OLEDB Client Provider. Running DBCC CHECKDB is a vital part of making sure your databases are healthy. Avoiding consistency checks just to lower OLEDB Wait Times is bad practice, and I strongly advise against it. Seeing high OLEDB Wait Times occur outside DBCC commands can indicate there is a performance issue somewhere in your SQL Server environment. If you are dealing with remote sources, such as linked servers or Excel files, you are also affected by the performance of the remote source. For instance, if you are querying information from a linked server and the linked server is experiencing performance problems, it will probably also be reflected in the OLEDB Wait Time. Also, certain operations, like sorts, can also impact the query duration on the linked server. Network connections to the remote source can also play a role in higher-than-normal

OLEDB Wait Times. If the network connection through which you are accessing your remote source experiences performance degradation, you will again notice this in the OLEDB Wait Times.

Because the OLEDB wait can occur for various reasons, some of which are benign like DBCC commands, and some that can be related to performance issues, I advise you not to ignore the OLEDB Wait Type, but rather to monitor it like other performance-indicating Wait Types.

OLEDB Summary

The OLEDB Wait Type can occur due to various sources that use the Object Linking and Embedding Database (OLEDB) Client Provider. Most of the sources are related to remote data sources, like linked servers. Some internal processes also use the OLEDB Client Provider, most notably the DBCC command. Seeing higher-than-normal Wait Times on the OLEDB Wait Type doesn't have to mean there is a performance problem, especially when they can be correlated to a planned DBCC command execution. Seeing higher-than-normal Wait Times outside DBCC command when you are using remote data sources like linked servers can mean that the remote data source is experiencing performance problems. In this case, focus on the data source; if the source has problems it is bound to affect the OLEDB Wait Times as well.

TRACEWRITE

The TRACEWRITE Wait Type is a special Wait Type that only collects Wait Time when a trace is running, and most commonly a SQL Profiler trace. A trace is a background process in SQL Server that collects various, often user-specified, information about the performance of a SQL Server instance. For example, it is possible to use SQL Server Profiler to capture currently executing queries, filtered against a single database, with runtime information. There are various trace methods available in SQL Server, but the most common one that affects the TRACEWRITE Wait Type is the SQL Server Profiler Trace.

SQL Server Profiler is an application that is part of the SQL Server suite and allows users to create and monitor traces against SQL Server instances. The SQL Server Profiler was announced as deprecated by Microsoft with the introduction of SQL Server 2012, and Microsoft recommends you use Extended Events to capture traces. Even though the SQL Server Profiler is deprecated, it is still available in SQL Server 2014, and many people still rely on SQL Server Profiler traces instead of Extended Events to troubleshoot and monitor query performance.

The bad news about using SQL Server Profiler is that it can cause some performance overhead while a trace is being performed. Microsoft released an article that concluded that running SQL Server Profiler traces on busy systems can have an impact of 10% on the amount of transactions per seconds; you can find the article here: <https://msdn.microsoft.com/en-us/library/cc293614.aspx>. Because SQL Server Profiler traces can have such a big impact on the performance of your system, I believe it is important to monitor the TRACEWRITE Wait Time.

What Is the TRACEWRITE Wait Type?

As we just noted, the TRACEWRITE Wait Type will show up on your system when traces are being performed against your SQL Server instance using the SQL Server Profiler. Because SQL Server Profiler traces can have such an impact on the performance of your SQL Server instance, it is advisable that you monitor the Wait Type to detect if any SQL Server Profiler traces are being performed.

There are a variety of reasons why you would want to run a SQL Server Profiler trace; for instance, if you want to troubleshoot a very specific query problem or when monitoring how many times a specific query gets executed. Even though there are alternatives to the SQL Server Profiler, like server-side traces and Extended Events, the SQL Server Profiler tool is very easy to use compared to the often complex Extended Events.

TRACEWRITE Example

To show you an example of TRACEWRITE waits occurring, we are going to have to start a SQL Server Profiler trace. The SQL Server Profiler program is part of the Management Tools - Complete feature, which you can select when installing SQL Server, or when adding features to an existing installation. Figure 11-9 shows the feature inside the SQL Server Setup.

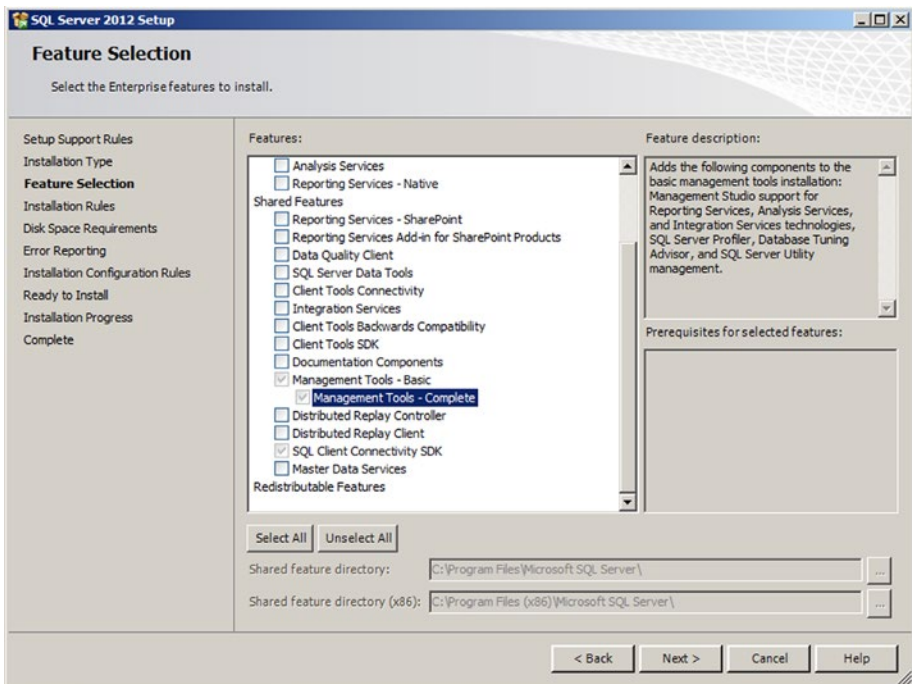


Figure 11-9. Management Tools - Complete feature in SQL Server Setup

When you have installed the Management Tools - Complete feature you can find the SQL Server Profiler in the SQL Server - Performance Tools folder underneath the Start menu, or in the C:\Program Files (x86)\Microsoft SQL Server\[edition number]\Tools\Binn folder.

After starting up the SQL Server Profiler you can start a new trace by clicking the New Trace button, shown in Figure 11-10, or selecting File ► New Trace.




Figure 11-10. New SQL Server Profiler trace

When you start a new trace you will need to connect to the SQL Server instance you want to trace. In this case I connected to my test SQL Server instance. After logging on to the SQL Server instance, the Trace Properties window will open. This window will give you a variety of options with which to configure your trace and how you want to store your trace. In this example we are not going to change anything in the General tab, but instead will go directly to the Events Selection tab. There we can select what events we want to capture, and optionally supply filters for those events. By default a selection is preloaded when starting a new trace, as you can see in Figure 11-11.

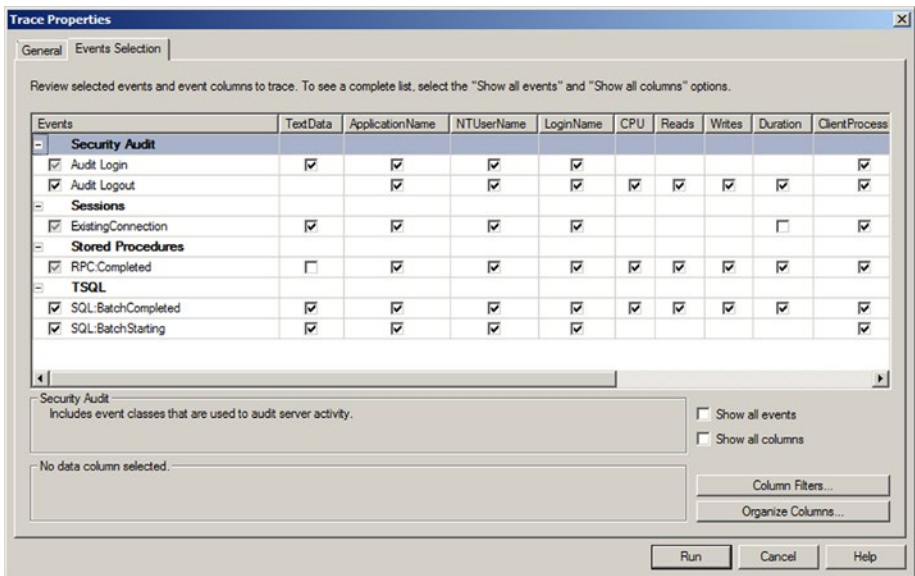


Figure 11-11. SQL Server Profiler default event selection

For this example we do not need all the extra events that are selected by default. In this case we check the Show all columns checkbox, and only select the SQL:BatchCompleted event, as you can see in Figure 11-12. This will record all the T-SQL statements that are executed against the test SQL Server instance and capture all the information available for the event.

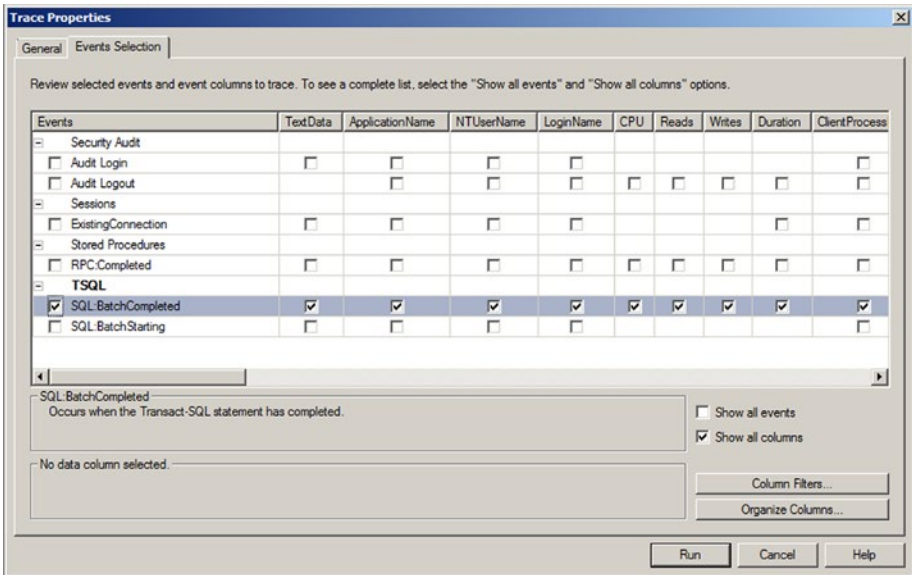


Figure 11-12. SQL:BatchCompleted event selected

We won't configure any filters on the event, so we will capture every T-SQL statement we execute against the SQL Server instance. We press Run to start the trace, which will open the trace window that will show us the events when they take place on our SQL Server instance along with additional information about, in this case, the query.

Now that our SQL Server Profiler trace is running, we should be able to notice TRACEWRITE waits occurring. We execute the query that follows in SQL Server Management Studio against the sys.dm_os_waiting_tasks DMV:

```
SELECT *
FROM sys.dm_os_waiting_tasks
WHERE wait_type = 'TRACEWRITE';
```

The results of this query are shown in Figure 11-13:

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address
1	0x000000017F02DC38	57	0	581	TRACEWRITE	0x0000000000000001

Figure 11-13. TRACEWRITE waits

Even though we are not running any workload on the test SQL Server instance, the TRACEWRITE Wait Type will still be logged. This is normal since the TRACEWRITE Wait Type will always be recorded as long as a SQL Server Profiler trace is active.

Lowering TRACEWRITE Waits

As I mentioned before, if you notice TRACEWRITE waits occurring it means someone is running a SQL Server Profiler trace against your SQL Server instance. Because a SQL Server Profiler trace can have such a big impact on the performance of your SQL Server instance, it is important to know who is running the SQL Server Profiler trace and why.

Thankfully, there is a catalog view we can query to view trace activity—the `sys.traces` view. The `sys.traces` catalog view will give you an overview of traces that are either active or paused against the SQL Server instance. The query that follows will retrieve all the information inside the `sys.traces` catalog view:

```
SELECT *
FROM sys.traces;
```

Running this query against the test SQL Server instance returns the information shown in Figure 11-14 (some columns did not fit inside the image).

	id	status	path	max_size	stop_time	max_files	is_rowset	is_rollover	is_shutdown	is_default
1	1	1	C:\Program Files\Microsoft SQL Server\MSSQL11.MSS...	20	NULL	5	0	1	0	1
2	2	1	NULL	NULL	NULL	NULL	1	0	0	0

Figure 11-14. `sys.traces`

Some important columns I want to highlight from the `sys.traces` catalog view are the `status` and `reader_spid` columns. The `status` column returns either a 0 or a 1, where a 0 indicates the trace is stopped or paused and a 1 indicates the trace is currently running. The `reader_spid` column returns the session ID of the session that started the trace. We can use this information to detect who is running the trace.

In our case, the trace we started in the example has an ID of 2, while the ID of 1 is reserved for the background SQL Server trace that is, by default, always active. This default trace collects specific information about the health of the SQL Server instance and can be used when troubleshooting. Because it is a so-called server-side trace, it does not record TRACEWRITE Wait Time while it is running.

Now that you can identify the user that is running the trace you can take action if you believe the trace has a negative effect on the performance of your SQL Server instance.

After stopping SQL Server Profiler traces in order to lower the TRACEWRITE Wait Time, there are other methods available if you really need to capture traces against your SQL Server instance. The most logical one is recreating your SQL Server Profiler trace within an Extended Event session. Extended Events have a much smaller overhead than SQL Server Profiler traces and allow even more events and options while capturing traces.

If you still want to use SQL Server Profiler to analyze traces, it can be a good idea to convert the trace you would normally run in the SQL Server Profiler application to a server-side trace. Just like with Extended Events, server-side traces have minimal overhead compared to traces that are performed through the SQL Server Profiler application. Let's convert the SQL Server Profiler trace we created in the example section to a server-side trace and monitor the effects on the TRACEWRITE Wait Type.

The easiest way to convert a SQL Server Profiler trace is by defining the trace in the SQL Server Profiler application without starting it. Instead, select the File ► Export ► Script Trace Definition ► For SQL Server 2005 - SQL11 option, as shown in Figure 11-15.

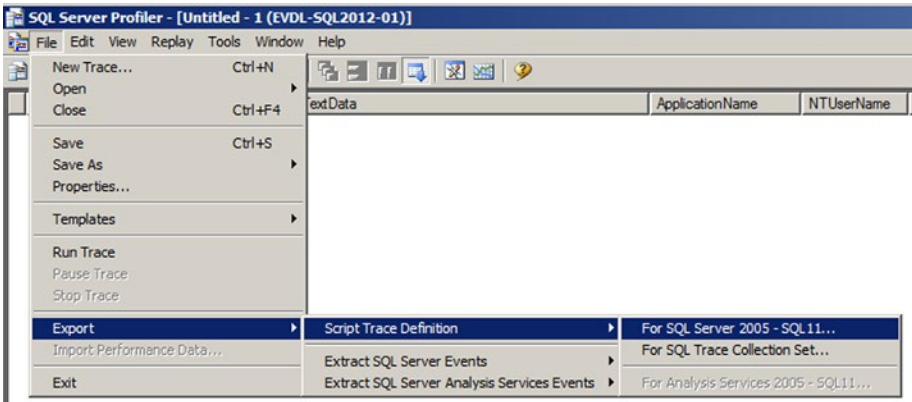


Figure 11-15. Export SQL Server Profiler trace to trace definition

After clicking the File ► Export ► Script Trace Definition ► For SQL Server 2005 - SQL11 option we will be asked to save a .sql file. The entire trace definition will be scripted inside this .sql file. We can open this file in SQL Server Management Studio, modify the file location and some other options inside the script, and execute it. This will return the ID of the trace we just created and save the trace information to a file we specified at the top of the script. Figure 11-16 shows a part of the exported trace definition on our test SQL Server instance.


```

1  /*****
2  /* Created by: SQL Server 2012 Profiler          */
3  /* Date: 06/18/2015  04:02:59 PM              */
4  /*****/
5
6
7  -- Create a Queue
8  declare @rc int
9  declare @TraceID int
10 declare @maxfilesize bigint
11 set @maxfilesize = 5
12
13 -- Please replace the text InsertFileNameHere, with an appropriate
14 -- filename prefixed by a path, e.g., c:\MyFolder\MyTrace. The .trc extension
15 -- will be appended to the filename automatically. If you are writing from
16 -- remote server to local drive, please use UNC path and make sure server has
17 -- write access to your network share
18
19 exec @rc = sp_trace_create @TraceID output, 0, N'C:\temp', @maxfilesize, NULL
20 if (@rc != 0) goto error
21
22 -- Client side File and Table cannot be scripted
23
24 -- Set the events
25 declare @on bit
26 set @on = 1
27 exec sp_trace_setevent @TraceID, 12, 1, @on
28 exec sp_trace_setevent @TraceID, 12, 9, @on
29 exec sp_trace_setevent @TraceID, 12, 3, @on

```

Figure 11-16. Trace definition

After executing the script to create a server-side trace, we received a trace ID of 2. The trace ID is very important because it is the only way to either start or stop the server-side trace. After creation, the server-side trace is automatically started. If we query the `sys.traces` catalog view, we can see the server-side trace that was just created, as shown in Figure 11-17.

	id	status	path	max_size	stop_time	max_files	is_rowset	is_rollover	is_shutdown	is_default
1	1	1	C:\Program Files\Microsoft SQL Server\MSSQL11.MSS...	20	NULL	5	0	1	0	1
2	2	1	C:\temp\trace_file.trc	5	NULL	1	0	0	0	0

Figure 11-17. `sys.traces`

The only way to interact with the server-side trace we created is to execute the `sp_trace_setstatus` stored procedure and supply the trace ID and a status ID. For instance, executing the query that follows will stop the server-side trace with a trace ID of 2:

```
EXEC sp_trace_setstatus 2, 0
```

To start it again we can execute this command:

```
EXEC sp_trace_setstatus 2, 1
```

And finally, to close the trace entirely we can execute the following command:

```
EXEC sp_trace_setstatus 2, 3
```

This does not delete the server-side trace though. As a matter of fact, server-side traces are only removed by a restart of the SQL Server service.

Because a server-side trace can only capture to a trace file, you can navigate to the file you supplied in the server-side trace definition and open the file in SQL Server Profiler. Thus, you can capture the same information as by using the SQL Server Profiler application but at a much lower performance price.

TRACEWRITE Summary

The TRACEWRITE Wait Type indicates a SQL Server Profiler trace is currently being performed against the SQL Server instance. SQL Server Profiler traces can have a pretty big impact on the performance of your SQL Server instance, and for this reason it is important to monitor the number of traces running against your SQL Server instance. Thankfully, there are some alternatives to SQL Server Profiler traces. You can either choose to convert your SQL Server Profiler trace to an Extended Events session or execute the SQL Server Profiler trace using server-side tracing.

WAITFOR

The final Wait Type in this chapter is one of the few Wait Types that are directly related to a T-SQL command. The WAITFOR Wait Type doesn't indicate performance problems, though it definitely has an impact on the duration of the query that is executing the related WAITFOR T-SQL command.

What Is the WAITFOR Wait Type?

The WAITFOR Wait Type will get recorded whenever a query is being executed that uses the WAITFOR command. The WAITFOR T-SQL command will stop the execution of the query until a specific amount of time has passed or a specific point in time has been reached. When that happens, the query execution will continue. The WAITFOR command is frequently used inside queries or scripts to force a pause inside the query execution. For instance, in Chapter 3, “Building a Solid Baseline,” we used the WAITFOR command to wait a specific amount of time so we could compare two measurements taken 15 minutes apart.

While pausing the query execution using the `WAITFOR` command, the transaction holding the `WAITFOR` command will remain open until the entire transaction has completed. This means that threads are being held by the transaction that cannot be used for other processes. SQL Server also reserves a dedicated thread just for the `WAITFOR` command; if too many threads are associated with `WAITFOR` commands and thread starvation occurs, SQL Server will select random `WAITFOR` threads and terminate them to free up more threads.

In many cases the `WAITFOR` command is explicitly used by the person who wrote the query or script, and in that sense only impacts that specific query or script; thus, there is no reason to be alarmed when seeing high `WAITFOR` Wait Times occur. It just indicates that queries are using the `WAITFOR` command.

WAITFOR Example

To show you a quick example of the `WAITFOR` Wait Type, you can execute the query in Listing 11-4. The query will reset the `sys.dm_os_wait_stats` DMV, execute a `WAITFOR DELAY` statement that causes the script execution to wait for 30 seconds, and then query the `sys.dm_os_wait_stats` DMV for `WAITFOR` waits.

Listing 11-4. `WAITFOR` waits

```
DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR);

WAITFOR DELAY '00:00:30';

SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'WAITFOR';
```

When the query in Listing 11-4 finishes, you should see that one `WAITFOR` wait occurred, having a total Wait Time of roughly 30 seconds, as you can see in Figure 11-18.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	WAITFOR	1	30007	30007	0

Figure 11-18. `WAITFOR` wait

WAITFOR Summary

The `WAITFOR` Wait Type is one of the few Wait Types that are directly related to the execution of a T-SQL command, in this case `WAITFOR`. The `WAITFOR` Wait Type doesn't indicate any performance problems with your SQL Server instance, it just indicates the `WAITFOR` command is being used by a query or script. The `WAITFOR` T-SQL command will only impact the execution time of the query or script that uses it; therefore, the only way to lower `WAITFOR` Wait Times is by removing the `WAITFOR` command inside queries.



In-Memory OLTP–Related Wait Types

With the release of SQL Server 2014, Microsoft introduced a brand new SQL Server feature called In-Memory OLTP (or codename Hekaton). In-Memory OLTP is a memory-optimized database engine that is directly integrated into the SQL Server 2014 SQL Server Engine. In-Memory OLTP is an enterprise-only feature designed to improve performance—up to 20 times, according to Microsoft—by placing tables entirely into the memory of your SQL Server instance. These memory-optimized tables are fully durable and use lock-and-latch free structures to optimize concurrency control.

With the introduction of In-Memory OLTP, various new Wait Types have been added to SQL Server 2014. Most of these are recognizable by the `_XTP_` (or eXtreme Transaction Processing) section in the Wait Type name. In this chapter we will take a look at some of these new, In-Memory OLTP-related Wait Types in SQL Server 2014. At the time of writing this book, information about these new Wait Types is extremely limited. I believe one reason for the lack of information is that the use-cases for In-Memory OLTP are very limited. Personally, I have yet to come across a client who uses this new technology in production, but because Microsoft gives a lot of attention to In-Memory technologies, I believe the number of use-cases and production systems that use In-Memory OLTP will only increase in the future.

Because of the lack of information about In-Memory OLTP Wait Types, the analyses of the Wait Types discussed in this chapter are based on so-called educated guesses done after many hours of testing In-Memory OLTP behavior. This also changes the format of the sections discussing the specific Wait Types. I removed the “Lowering Wait Times” section, because there simply is too little information on which to base conclusive and solid advice on how to lower the Wait Times of In-Memory OLTP Wait Types. I also removed the “Example” section and merged it with the “Wait Type Definition” section. With these changes, you can see this chapter as a kind of journal I wrote during my In-Memory OLTP Wait Type feature analysis.

Before we dive into the Wait Types though, let’s first take a (simplified and short) look at what In-Memory OLTP is and how it works. I will focus on memory-optimized tables during this chapter. In-Memory OLTP also introduced other features, like natively-compiled stored procedures and hash indexes, but these are beyond the scope of this chapter.

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-1-4842-1139-7_12](https://doi.org/10.1007/978-1-4842-1139-7_12)) contains supplementary material, which is available to authorized users.

Introduction to In-Memory OLTP

The main difference between traditional, disk-based tables and memory-optimized tables is that memory-optimized tables reside completely in the memory of your SQL Server instance. Unlike traditional tables, where data pages from that table are moved from disk into memory and back out again, memory-optimized tables are moved to your system's memory at SQL Server startup and never leave the memory (unless the memory-optimized table is removed, of course). While this might sound a bit scary at first, memory-optimized tables are, by default, fully durable. This means that if your SQL Server instance crashes, memory-optimized table data is not lost. Of course, having an entire table reside in the memory of your SQL Server instance also has its disadvantages. You need to make sure you have enough free memory to accommodate the entire memory-optimized table (and some extra memory to accommodate row versions used when accessing such tables). Calculating the memory requirements can be difficult, but the following article can help you out <https://msdn.microsoft.com/en-us/library/dn282389.aspx>.

The memory you reserve for memory-optimized tables is claimed by SQL Server and will not be swapped out; if your memory-optimized tables use too much memory, your SQL Server instance will run into memory starvation issues that cause performance degradation or, worst-case, cause SQL Server to crash. This is a major difference compared to, for instance, the buffer cache, where pages are swapped out of memory when memory pressure occurs. Another disadvantage is that many data types or SQL Server features are not supported for memory-optimized tables. The complete list of what can and cannot be used can be found here: [https://msdn.microsoft.com/en-us/library/dn246937\(v=sql.120\).aspx](https://msdn.microsoft.com/en-us/library/dn246937(v=sql.120).aspx) and here: [https://msdn.microsoft.com/en-us/library/dn133181\(v=sql.120\).aspx](https://msdn.microsoft.com/en-us/library/dn133181(v=sql.120).aspx). While writing this chapter, Microsoft released SQL Server 2016 CTP2, and in this build some major limitations were resolved. I expect that with future releases more and more limitations will be removed, making memory-optimized tables more attractive.

So, how do memory-optimized tables work, and why do they perform that much faster than traditional disk-based tables? Let's take a look at some of the internals of In-Memory OLTP.

CFPs

Like I mentioned earlier, by default memory-optimized tables are durable (you can choose to create a non-durable table that has its contents cleared on a SQL Server Service restart, but you have to explicitly specify this). The way this durability is achieved is through so-called checkpoint file pairs (CFPs). CFPs consist of two files, a data and a delta file, that exist inside a special memory-optimized filegroup that you have to create for the database where you want to use memory-optimized tables.

Unlike traditional tables that store row data inside data pages, data files store the rows of *all* your memory-optimized tables. I emphasize the word *all* because a single data file can hold the rows of many memory-optimized tables, unlike data pages that store row data for a single table. The rows inside a data file are stored sequentially based on the time they were inserted into a memory-optimized table. This is different than with data pages, which hold row information inside extents for traditional tables. Because rows are stored sequentially inside the data files, there is a performance increase when reading rows since it eliminates the random reads that occur when reading rows from traditional tables. Figure 12-1 shows an abstract view of a data file and the row data it holds.




Figure 12-1. Memory-optimized table's data file

There are always multiple data files inside a memory-optimized filegroup. When you first create the memory-optimized filegroup, SQL Server will automatically pre-allocate a number of data files in the file location of the memory-optimized filegroup. The data files will always have a fixed file size, either 128 MB on systems with more than 16 GB memory, or 16 MB when there is less than or equal to 16 GB memory. When a data file is full, a new data file will automatically be created and new rows will be inserted into the new data file. It is important to know that the data file keeps track of rows based on the transaction-commit timestamp that inserted the row into the data file (shown by the number inside Figure 12-1). Even if new data files are added and rows are spread across multiple data files, the data files will always have a contiguous range of transactions. Figure 12-2 shows multiple data files and the transaction-commit timestamps associated with those data files.

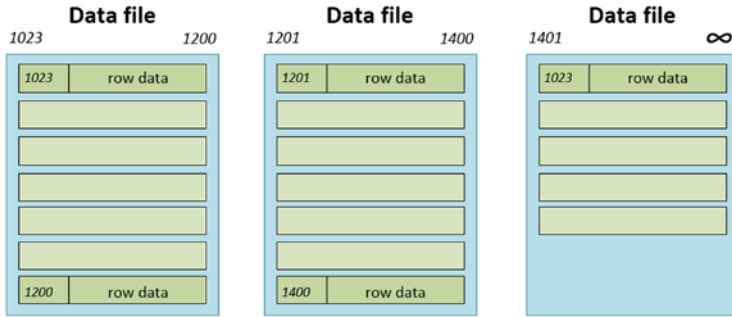


Figure 12-2. Data files and transaction timestamps

Notice in Figure 12-2 that the last data file still has room for new rows—it doesn't have a transaction timestamp to indicate the file is full, so new rows will be added to that file.

Another important characteristic of the data file is that rows that are deleted are not directly removed from the file. Instead they are tracked by the delta file that is associated with the data file. The delta file logs any deletes made in the data file and is connected to the data file by the transaction timestamp range. Row updates for memory-optimized tables are tracked as a delete and insert operation.

The population of the data and delta files is performed by a background thread—called the offline checkpoint thread—that runs constantly in the background of SQL Server. This is different than the checkpoint process used for traditional tables where pages are written to the database data files at intervals. The offline checkpoint thread monitors the transaction log for operations performed on memory-optimized tables and directly writes to the data and delta files.

Over time, when data files accumulate more deleted rows, a merge operation will take place that will merge multiple data files together into one data file. The merge operation will create new data and delta files and move the contents of one or more data and delta files into the new files, but it will not move the rows that were marked as deleted. The transaction-commit timestamps will be adjusted in the new data and delta files so they match the timestamps of the files that were merged. Figure 12-3 shows a simplified view of a merge operation on a data-file level. Keep in mind that a merge will also impact the delta file.

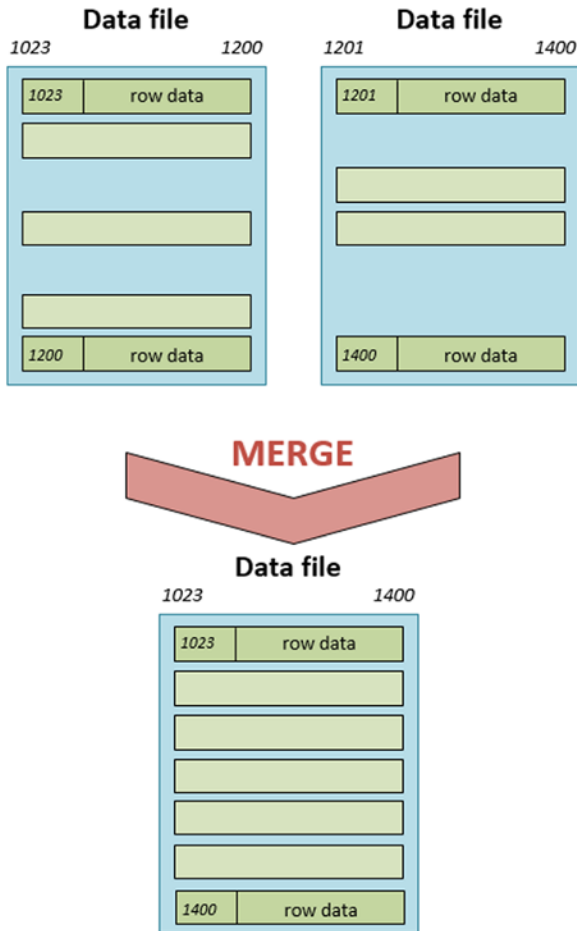


Figure 12-3. Merge operation

Isolation

Concurrent access to memory-optimized tables is handled through snapshot-based transaction isolation. This isolation level shares many characteristics with the snapshot isolation we can use on disk-based tables, but there are some differences so as to optimize throughput. First of all, the snapshot-based transaction isolation uses row versions when concurrent transactions want to access the same row. Instead of storing the row versions in the TempDB database like regular snapshot isolation, the row versions for memory-optimized tables are stored in-line in the data files itself.

Another difference is that snapshot-based transaction isolation uses an optimistic concurrency control. This means that SQL Server assumes no transaction conflict will occur when concurrent transactions access the same data. Because of this assumption

there is no need for locks or latches to protect the memory-optimized table data. There is a form of conflict detection active, however, and when it detects that a conflict has occurred it will end one of the transactions, and that transaction will need to be retried.

Not having to place and maintain locks and latches is another major contribution to the performance of In-Memory OLTP.

Transaction Log Changes

The final differences I want to discuss in this section are the modifications to the behavior of the transaction log regarding memory-optimized tables. For traditional tables, a log record will be generated when a transaction starts whether it gets committed or not. For memory-optimized tables, the log record will only be generated when the transaction begins the commit processing. This means no information for transactions that are rolled back is recorded. This minimizes interaction with the transaction log on disk, thus improving performance.

Another modification is that changes to indexes on memory-optimized tables are not logged in the transaction log. Since indexes that are created on memory-optimized tables are also maintained entirely in-memory, there is no need to record changes. Indexes on memory-optimized tables are regenerated on the start of the SQL Server Service.

The final difference I want to mention is the grouping of multiple transactions into one log record. For traditional tables every transaction will result in at least one log record. Transactions against memory-optimized tables are grouped together and then written as one log record (with a current maximum size of 24 KB). For instance, if you have 200 inserts against a traditional table, at least 200 log records would be generated. If we could fit 100 inserts into one log record for the memory-optimized table, we would only have two log records instead of at least 200. Again, this improves throughput for memory-optimized tables.

Now that we have taken a (simplified and short) look at some of the inner workings of memory-optimized tables, let's move on to some of the Wait Types that are related to In-Memory OLTP. Most of the three Wait Types that we will discuss in this chapter are related, one way or another, to the new offline checkpoint process introduced with In-Memory OLTP.

WAIT_XTP_HOST_WAIT

The first Wait Type we will discuss in this chapter is `WAIT_XTP_HOST_WAIT`. This Wait Type shares some characteristics with the `CHECKPOINT_QUEUE` Wait Type we discussed in Chapter 11, "Background and Miscellaneous Wait Types," in that it seems to be running continuously but only writes its wait information to `sys.dm_os_wait_stats` at specific conditions.

What Is the `WAIT_XTP_HOST_WAIT` Wait Type?

If we look up some information about the `WAIT_XTP_HOST_WAIT` Wait Type on Books Online, we get a not-so-helpful definition: "Occurs when waits are triggered by the database engine and implemented by the host." This doesn't give us a lot of clues about the processes that might be related to the `WAIT_XTP_HOST_WAIT` Wait Type, which means we have to do a little bit of digging ourselves.

Before we can start investigating the `WAIT_XTP_HOST_WAIT` Wait Type we need to create a memory-optimized table. I used the script shown in Listing 12-1 to create a new database with a single memory-optimized table. There are some path references in this script that you will need to change to make sure the database data and log files are created in the right location.

Listing 12-1. Create test database and memory-optimized table

```
-- Create database
-- Make sure to change the file locations if needed
USE [master]
GO

CREATE DATABASE [OLTP_Test] CONTAINMENT = NONE
ON PRIMARY
(
    NAME = N'OLTP_Test', FILENAME = N'E:\Data\OLTP_Test_Data.mdf' ,
    SIZE = 51200KB , FILEGROWTH = 10%
)
LOG ON
(
    NAME = N'OLTP_Test_log', FILENAME = N'E:\Log\OLTP_Test_Log.ldf' ,
    SIZE = 10240KB , FILEGROWTH = 10%
);
GO

-- Add the Memory-Optimized Filegroup
ALTER DATABASE OLTP_Test ADD FILEGROUP OLTP_MO CONTAINS MEMORY_OPTIMIZED_DATA;
GO

-- Add a file to the newly created Filegroup.
-- Change drive/folder location if needed.
ALTER DATABASE OLTP_Test ADD FILE (name='OLTP_mo_01', filename=
'E:\data\OLTP_Test_mo_01.ndf') TO FILEGROUP OLTP_MO;
GO

-- Create our test table
USE [OLTP_Test]
GO

CREATE TABLE OLTP
(
    ID INT IDENTITY (1,1) PRIMARY KEY NONCLUSTERED,
    RandomData1 VARCHAR(50),
    RandomData2 VARCHAR(50),
    ID2 UNIQUEIDENTIFIER
)
WITH (MEMORY_OPTIMIZED=ON);
GO
```

Now that we have a memory-optimized table we can use for testing, let’s take a look at the `WAIT_XTP_HOST_WAIT` Wait Type in the `sys.dm_os_wait_stats` and `sys.dm_os_waiting_tasks` DMV using the following query:

```
SELECT *
FROM sys.dm_os_waiting_tasks
WHERE wait_type = 'WAIT_XTP_HOST_WAIT';

SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'WAIT_XTP_HOST_WAIT';
```

The results of this query on my test SQL Server instance can be seen in Figure 12-4.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address
1	0x00000001FBA068C8	10	0	168588	WAIT_XTP_HOST_WAIT	NULL

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	WAIT_XTP_HOST_WAIT	4	121051	118419	3

Figure 12-4. `WAIT_XTP_HOST_WAIT` waits

The first thing you’ll notice is that the `WAIT_XTP_HOST_WAIT` Wait Type is constantly showing up in the `sys.dm_os_waiting_tasks` DMV, and over time the Wait Time increases. Also, the `session_id` that is related to the `WAIT_XTP_HOST_WAIT` Wait Type indicates it is an internal SQL Server thread that is recording the wait. From this information we can already formulate some conclusions about the `WAIT_XTP_HOST_WAIT` Wait Type—it is related to an internal background process that continuously runs. What’s also interesting is that if you were to run the query a second time, the Wait Time in the `sys.dm_os_waiting_tasks` DMV increases but the Wait Time in the `sys.dm_os_wait_stats` remains the same. So far we have run into one other Wait Type that shared this characteristic, the `CHECKPOINT_QUEUE` Wait Time, which we discussed in Chapter 11, “Background and Miscellaneous Wait Types.”

Since the `CHECKPOINT_QUEUE` Wait Type has interesting behavior in that it only writes the accumulated Wait Time of the `sys.dm_os_waiting_tasks` DMV to the `sys.dm_os_wait_stats` DMV when an automatic checkpoint occurs, I decide to simply run a checkpoint command against the `OLTP_Test` database I created using Listing 12-1, and then query both the DMVs again. The impact on the Wait Times of the `WAIT_XTP_HOST_WAIT` Wait Type can be seen in Figure 12-5.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address
1	0x00000001FBA068C8	10	0	435	WAIT_XTP_HOST_WAIT	NULL

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	WAIT_XTP_HOST_WAIT	6	1066941	919386	4

Figure 12-5. `WAIT_XTP_HOST_WAIT` wait information after checkpoint

As you can see in Figure 12-5, the Wait Time in the `sys.dm_os_waiting_tasks` DMV is very small again, but the Wait Time in the `sys.dm_os_wait_stats` DMV has increased a lot. Because of this behavior, I believe the `WAIT_XTP_HOST_WAIT` Wait Type has something to do with the offline checkpoint process that is related to memory-optimized tables.

To verify my guess I need to dig a little bit deeper to find out what goes on underneath the hood of SQL Server when a checkpoint is performed against a memory-optimized table. To get this information I create an Extended Events session that captures the call stack whenever SQL Server runs into a `WAIT_XTP_HOST_WAIT` wait. I won't bore you with the methods I use for creating this Extended Events session here, but Paul Randal wrote an amazing blog post about capturing call stacks whenever a specific wait occurs that you can use to collect some call stacks yourself. You can find Paul's blog post here: <http://www.sqlskills.com/blogs/paul/determine-causes-particular-wait-type/>.

The results of my Extended Events session for capturing the call stack when a `WAIT_XTP_HOST_WAIT` wait occurs is shown here:

```
sqlldk.dll!XeSosPkg::wait_info::Publish+0x138
sqlldk.dll!SOS_Scheduler::UpdateWaitTimeStats+0x26c
sqlldk.dll!SOS_Task::PostWait+0x9e
sqlmin.dll!EventInternal&lt;SuspendQueueSlock&gt;;:Wait+0x1fb
sqlmin.dll!HkHostWait::Wait+0xce
hkengine.dll!CkptFilePair::CreateInstance+0x61b
sqlmin.dll!HkHostReportFailure::KillProcess+0x372
sqlldk.dll!SOS_Task::Param::Execute+0x21e
sqlldk.dll!SOS_Scheduler::RunTask+0xa8
sqlldk.dll!SOS_Scheduler::ProcessTasks+0x279
sqlldk.dll!SchedulerManager::WorkerEntryPoint+0x24c
sqlldk.dll!SystemThread::RunWorker+0x8f
sqlldk.dll!SystemThreadDispatcher::ProcessWorker+0x3ab
sqlldk.dll!SchedulerManager::ThreadEntryPoint+0x226
kernel32.dll!BaseThreadInitThunk+0xd
ntdll.dll!RtlUserThreadStart+0x21
```

The first part I found really interesting is the inclusion of a new .dll file, `hkengine.dll`. Since In-Memory OLTP's codename was Hekaton, I am guessing that this .dll holds the new In-Memory OLTP functions, so let's zoom in on that particular call:

```
hkengine.dll!CkptFilePair::CreateInstance+0x61b
```

Seeing the function name, I am guessing it is related to checkpoint file pairs, and the `CreateInstance` bit suggests a new CFP was created when I executed the `CHECKPOINT` command. We can verify this by going to the location of the In-Memory file that we created in Listing 12-1. The interesting thing about adding a file to an In-Memory filegroup is that it will actually create a directory, and inside this directory there is a folder with a unique ID string. If you go further down the directory tree, you will end up in a folder with numbered files. Figure 12-6 shows a part of the contents of this folder on my test machine.

Name ^	Date modified	Type	Size
00000020-000000b0-0002	6/26/2015 8:50 AM	File	1,024 KB
00000020-000000b5-0003	6/26/2015 8:50 AM	File	16,384 KB
00000020-000000c0-0002	6/26/2015 8:50 AM	File	1,024 KB
00000020-000000c5-0003	6/26/2015 8:50 AM	File	16,384 KB
00000020-000000cc-0002	6/26/2015 8:50 AM	File	1,024 KB
00000020-000000d1-0003	6/26/2015 8:50 AM	File	16,384 KB
00000020-000000d8-0002	6/26/2015 8:50 AM	File	1,024 KB
00000020-000000dd-0003	6/26/2015 8:50 AM	File	16,384 KB
00000020-000000e4-0002	6/26/2015 8:50 AM	File	1,024 KB
00000020-000000e9-0003	6/26/2015 8:50 AM	File	16,384 KB
00000020-000000f0-0002	6/26/2015 8:50 AM	File	1,024 KB
00000020-000000f5-0003	6/26/2015 8:50 AM	File	16,384 KB
00000020-000000fc-0002	6/26/2015 8:50 AM	File	1,024 KB
00000020-000001b0-0002	6/26/2015 9:05 AM	File	16,384 KB
00000020-000001ba-0002	6/26/2015 9:05 AM	File	1,024 KB
00000020-000001bf-0003	6/26/2015 9:06 AM	File	16,384 KB
00000020-000001c6-0002	6/26/2015 9:05 AM	File	1,024 KB

Figure 12-6. *In-Memory filegroup files*

As a matter of fact, the files you are seeing here are the data and delta files that are associated with the memory-optimized table we created earlier. The 1 MB files are the delta files and the 16 MB ones are the data files.

Since I am guessing a checkpoint would create another CFP, I checked the number of files in the folder before executing a CHECKPOINT, which was 28 files. I then executed a CHECKPOINT command and looked at the number of files again, and it turned out there were now 30 files after the checkpoint.

WAIT_XTP_HOST_WAIT Summary

I believe the `WAIT_XTP_HOST_WAIT` Wait Type has a clear relation to the creation of new checkpoint file pairs. Apparently running a manual CHECKPOINT statement will generate a new CFP for the memory-optimized tables. Because the `WAIT_XTP_HOST_WAIT` Wait Type generates Wait Time constantly in the background, and writes it to the `sys.dm_os_wait_stats` DMV a new CFP was created (either by manual checkpoint, when an existing CFP was full, or when a Merge operation occurred), I believe the `WAIT_XTP_HOST_WAIT` Wait Type does not directly indicate performance problems. It mostly indicates that a new CFP has been added to the In-Memory filegroup. This does not mean this is the only process that generates `WAIT_XTP_HOST_WAIT` waits, though. There can be other processes that can also cause the waits, but so far they only occurred whenever a new CFP needed to be added.

WAIT_XTP_CKPT_CLOSE

The `WAIT_XTP_CKPT_CLOSE` Wait Type is another new Wait Type introduced in SQL Server 2014. As the name suggests, it seems to be related to the new offline checkpoint process introduced with the In-Memory OLTP feature.

What Is the `WAIT_XTP_CKPT_CLOSE` Wait Type?

The `WAIT_XTP_CKPT_CLOSE` Wait Type seems to be related to the new offline checkpoint process that was introduced in SQL Server 2014 with the release of In-Memory OLTP. As far as I can tell by analyzing the behavior of this Wait Type, it only records Wait Time when a checkpoint occurs, no matter if it is an automatic or manual checkpoint. The Wait Time the `WAIT_XTP_CKPT_CLOSE` Wait Type represents seems to be the time it takes for the checkpoint operation to complete. We can verify this easily by executing a `CHECKPOINT` command against the database and table we created earlier when we discussed the `WAIT_XTP_HOST_WAIT` Wait Type. I used the script in Listing 12-2 to clear the `sys.dm_os_wait_stats` DMV, insert a few rows inside the memory-optimized table, perform a `CHECKPOINT` operation, and then query the `sys.dm_os_wait_stats` DMV for `WAIT_XTP_CKPT_CLOSE` Wait Type information.

Listing 12-2. Generate `WAIT_XTP_CKPT_CLOSE` waits

```
USE [OLTP_Test];
GO

-- Clear sys.dm_os_wait_stats
DBCC SQLPERF('sys.dm_os_wait_stats', CLEAR);

-- Insert some rows
INSERT INTO OLTP
    (
        RandomData1,
        RandomData2,
        ID2
    )
VALUES
    (
        CONVERT(VARCHAR(50), NEWID()),
        CONVERT(VARCHAR(50), NEWID()),
        NEWID()
    );
GO 1000

-- Perform a CHECKPOINT
CHECKPOINT
```

```
-- Query sys.dm_os_wait_stats for WAIT_XTP_CKPT_CLOSE waits
SELECT *
FROM sys.dm_os_wait_stats
WHERE wait_type = 'WAIT_XTP_CKPT_CLOSE';
```

The results can be seen in Figure 12-7.

	wait_type	waiting_tasks_count	wait_time_ms	max_wait_time_ms	signal_wait_time_ms
1	WAIT_XTP_CKPT_CLOSE	2	1015	1014	0

Figure 12-7. WAIT_XTP_CKPT_CLOSE waits

Just as I did for the WAIT_XTP_HOST_WAIT Wait Type, I also captured the call stack when a WAIT_XTP_CKPT_CLOSE wait occurred:

```
sqldk.dll!XeSosPkg::wait_info::Publish+0x138
sqldk.dll!SOS_Scheduler::UpdateWaitTimeStats+0x2bc
sqldk.dll!SOS_Task::PostWait+0x9e
sqlmin.dll!EventInternal&lt;SuspendQueueLock&gt;::Wait+0x1fb
sqlmin.dll!HkCheckpointCtxtImpl::WaitForCkptComplete+0xd0
sqlmin.dll!HkHostWaitForCkptComplete+0x13a
sqlmin.dll!CheckpointWithOptionalTruncate+0xe6
sqllang.dll!CStmtCheckpoint::XretExecute+0xe7
sqllang.dll!CMsgExecContext::ExecuteStmts&lt;1,1&gt;+0x427
sqllang.dll!CMsgExecContext::FExecute+0xa33
sqllang.dll!CSource::Execute+0x86c
sqllang.dll!process_request+0xa57
sqllang.dll!process_commands+0x4a3
sqldk.dll!SOS_Task::Param::Execute+0x21e
sqldk.dll!SOS_Scheduler::RunTask+0xa8
sqldk.dll!SOS_Scheduler::ProcessTasks+0x279
sqldk.dll!SchedulerManager::WorkerEntryPoint+0x24c
sqldk.dll!SystemThread::RunWorker+0x8f
sqldk.dll!SystemThreadDispatcher::ProcessWorker+0x3ab
sqldk.dll!SchedulerManager::ThreadEntryPoint+0x226
kernel32.dll!BaseThreadInitThunk+0xd
ntdll.dll!RtlUserThreadStart+0x21
```

I believe the most interesting line is `sqlmin.dll!CheckpointWithOptionalTruncate+0xe6`, which seems to be the function that performs the truncate. It is followed by the `sqlmin.dll!HkCheckpointCtxtImpl::WaitForCkptComplete+0xd0` line that I believe records the time the previous checkpoint function took place, which gets posted later on to the Wait Statistics DMVs.

I don't believe seeing WAIT_XTP_CKPT_CLOSE waits occur is a direct cause for concern. They indicate that checkpoints are being performed. I can imagine that sudden high Wait Times for the WAIT_XTP_CKPT_CLOSE Wait Type can indicate a performance issue. As we saw in the previous section, performing a checkpoint against a memory-optimized table will

result in extra CFPs being created. I am guessing that if the allocation of CFPs takes a long time, the checkpoint operation will take longer to complete as well, resulting in higher `WAIT_XTP_CKPT_CLOSE` Wait Times. The amount of data a checkpoint has to process will probably also mean higher `WAIT_XTP_CKPT_CLOSE` Wait Times. Since the checkpoint writes data to the storage subsystem, the performance of your storage will probably also impact `WAIT_XTP_CKPT_CLOSE` Wait Times.

WAIT_XTP_CKPT_CLOSE Summary

The `WAIT_XTP_CKPT_CLOSE` Wait Type seems closely related to performing checkpoint operations. It indicates the time a checkpoint performed against a memory-optimized table took to complete. I don't believe this directly indicates performance issues, since it just records the time it took for the checkpoint to complete. The amount of work a checkpoint has to process will probably result in higher `WAIT_XTP_CKPT_CLOSE` Wait Times. Storage subsystem performance will probably also impact `WAIT_XTP_CKPT_CLOSE` Wait Times.

WAIT_XTP_OFFLINE_CKPT_NEW_LOG

The final In-Memory OLTP related Wait Type that I want to discuss in this chapter is the `WAIT_XTP_OFFLINE_CKPT_NEW_LOG` Wait Type. This is another Wait Type related to the offline checkpoint process that was introduced in SQL Server 2014.

What Is the WAIT_XTP_OFFLINE_CKPT_NEW_LOG Wait Type?

The `WAIT_XTP_OFFLINE_CKPT_NEW_LOG` Wait Type appears to be a benign Wait Type that records the length of time the offline checkpoint process is waiting for work. This is confirmed by Books Online, which has the following definition: "occurs when offline checkpoint is waiting for new log records to scan."

As we discussed earlier in this chapter, the offline checkpoint process monitors the transaction log for transactions that impact memory-optimized tables so those transactions can be recorded in the data and delta files. This is a constantly running process in the background of SQL Server, which means you will see an internal process with the `WAIT_XTP_OFFLINE_CKPT_NEW_LOG` Wait Type when you query the `sys.dm_os_waiting_tasks` DMV, as you can see in Figure 12-8.

	waiting_task_address	session_id	exec_context_id	wait_duration_ms	wait_type	resource_address	blocking_task_address
1	0x00000001f5b884e8	43	0	863	WAIT_XTP_OFFLINE_CKPT_NEW_LOG	NULL	NULL
2	0x00000001ff021088	28	0	3895	XE_TIMER_EVENT	NULL	NULL
3	0x00000001ff0168c8	25	0	100504	XE_DISPATCHER_WAIT	NULL	NULL
4	0x00000001f8f268c8	24	0	165	HADR_FILESTREAM_IOMGR_IOCOMPLETION	NULL	NULL
5	0x00000001f61104e8	23	0	990	SLEEP_TASK	NULL	NULL
6	0x00000001f6110108	22	0	45996534	BROKER_TRANSMITTER	NULL	NULL

Figure 12-8. `WAIT_XTP_OFFLINE_CKPT_NEW_LOG` waits inside `sys.dm_os_waiting_tasks`

Unlike the `WAIT_XTP_HOST_WAIT` Wait Type that only writes its Wait Time information to the `sys.dm_os_wait_stats` DMV when specific conditions occur, the `WAIT_XTP_OFFLINE_CKPT_NEW_LOG` Wait Type appears to wait for around five seconds, adds the Wait Time to the `sys.dm_os_wait_stats` DMV, and then resets the Wait Time in the `sys.dm_os_waiting_tasks` DMV again. This might suggest that the offline checkpoint process checks for new work at an interval of around five seconds.

To understand a little bit more about the offline checkpoint process, I captured a stack dump whenever a `WAIT_XTP_OFFLINE_CKPT_NEW_LOG` wait occurred. The stack dump gives us some interesting insight into the process itself, as you can see here:

```
sqlldk.dll!XeSosPkg::wait_info::Publish+0x138
sqlldk.dll!SOS_Scheduler::UpdateWaitTimeStats+0x2bc
sqlldk.dll!SOS_Task::PostWait+0x9e
sqlmin.dll!EventInternal&lt;SuspendQueueLock&gt;::Wait+0x1fb
sqlmin.dll!SequencedObject&lt;LogBlockId,SequencedWaitInfo&lt;LogBlockId&gt;,0&gt;::WaitUntilSequenceAdvances+0x160
sqlmin.dll!OfflineCheckpointWorker::GetNextLogBlock+0x10d
sqlmin.dll!OfflineCheckpointWorker::DoWorkInternal+0xf7
sqlmin.dll!OfflineCheckpointWorker::DoWork+0x3aa
sqlmin.dll!OfflineCheckpointWorker::WorkLoop+0x3fc
sqlldk.dll!SOS_Task::Param::Execute+0x21e
sqlldk.dll!SOS_Scheduler::RunTask+0xa8
sqlldk.dll!SOS_Scheduler::ProcessTasks+0x279
sqlldk.dll!SchedulerManager::WorkerEntryPoint+0x24c
sqlldk.dll!SystemThread::RunWorker+0x8f
sqlldk.dll!SystemThreadDispatcher::ProcessWorker+0x3ab
sqlldk.dll!SchedulerManager::ThreadEntryPoint+0x226
kernel32.dll!BaseThreadInitThunk+0xd
ntdll.dll!RtlUserThreadStart+0x21
```

The most interesting parts are when the `OfflineCheckpointWorker` function is being called. For readability, here is the section that involves the `OfflineCheckpointWorker` function:

```
sqlmin.dll!SequencedObject&lt;LogBlockId,SequencedWaitInfo&lt;LogBlockId&gt;,0&gt;::WaitUntilSequenceAdvances+0x160
sqlmin.dll!OfflineCheckpointWorker::GetNextLogBlock+0x10d
sqlmin.dll!OfflineCheckpointWorker::DoWorkInternal+0xf7
sqlmin.dll!OfflineCheckpointWorker::DoWork+0x3aa
sqlmin.dll!OfflineCheckpointWorker::WorkLoop+0x3fc
```

Seeing this stack dump makes me believe the offline checkpoint process is started, starts looking for work by reading log records from the transaction log (`LogBlock`), grabs the first `LogBlock` it needs to process, and loops until all `LogBlocks` are processed. When that is done, I suspect the offline checkpoint goes to sleep again, waits for around five seconds, then wakes up and checks for new log records.

Seeing this behavior makes me believe the `WAIT_XTP_OFFLINE_CKPT_NEW_LOG` Wait Type is harmless. It just indicates that the offline checkpoint process is waiting for work to arrive.

`WAIT_XTP_OFFLINE_CKPT_NEW_LOG` Summary

The `WAIT_XTP_OFFLINE_CKPT_NEW_LOG` Wait Type is related to the offline checkpoint process and indicates that process is waiting for work to arrive. Because the `WAIT_XTP_OFFLINE_CKPT_NEW_LOG` Wait Type only indicates that the offline checkpoint process is waiting for work, I believe the Wait Type doesn't indicate any performance issues and can probably be safely ignored.

All of the Wait Type analyses I did in this chapter are based on my own observations. Therefore I had to make some guesses on how certain Wait Types behave, based on those observations. Because there is so little information available about In-Memory OLTP Wait Types at the time of writing this book, and there wasn't material available to compare my observations against, this chapter results in theories that might not be conclusive. I believe that when In-Memory OLTP gets used more in production environments, information about the related Wait Types will become easier to find, leading to more definite explanations. I didn't want to withhold the information I gathered about these new Wait Types though, and by adding this chapter I hope to give you a sense of direction when dealing with In-Memory OLTP Wait Types.

APPENDIX I



Example SQL Server Machine Configurations

During the writing of this book I used a few different test systems to generate the examples that are used. This appendix will describe the configuration of the systems I used during the examples and Wait Type demonstrations. If I needed to modify the system to demonstrate a specific Wait Type or situation occurring, this will be included in the text inside the chapter that holds the demonstration.

All my test systems are virtual machines I created inside Oracle VirtualBox, a free-to-use virtualization software product that you can download from <https://www.virtualbox.org/>.

Another tool I frequently used during examples is Ostress. Ostress is part of the RML utilities provided to manage your SQL Server's performance. You can download the RML utilities using this link: <https://www.microsoft.com/en-us/download/details.aspx?id=4511>.

Default Test Machine

The table that follows shows the virtual machine configuration I used for the majority of the book, except for Chapter 9, which discusses high-availability and disaster-recovery Wait Types, and Chapter 12, which discusses In-Memory OLTP-related Wait Types.

Configuration	Value
Computer name	EVDL-SQL2012-01
vCPUs	4
Architecture	64-bit
Memory	2 GB
Storage	20 GB System Drive C:\ (SSD) 25 GB Data Drive D:\ (SSD)
Data Drive Layout	D:\Data, MDF Files D:\Log, LDF Files D:\Backup, Backup files

(continued)

Configuration	Value
Operating System	Windows Server 2008R2 SP1
SQL Server Edition	SQL Server 2012 SP1 Enterprise
SQL Server Features	Database Engine Services Management Tools - Complete
SQL Server Instance Name	MSSQLSERVER (Default instance)

HA/DR Test Machines

The tables that follow show the configurations of the virtual machines I used for demonstrating high-availability and disaster-recovery Wait Types as described in Chapter 9.

Configuration	Value
Computer name	EVDL-DC-01
Role	Domain Controller (PROWAITTS)
vCPUs	1
Architecture	64-bit
Memory	512 MB
Storage	20 GB System Drive C:\ (SSD)
Operating System	Windows Server 2008R2 SP1
Computer name	EVDL-SQL-AG01
Role	Principle (mirroring) Primary (AlwaysOn) Failover Cluster node
vCPUs	2
Architecture	64-bit
Memory	2 GB
Storage	25 GB System Drive C:\ (SSD) 20 GB Data Drive D:\ (SSD)
Data Drive Layout	D:\Data, MDF Files D:\Log, LDF Files D:\Backup, Backup files
Operating System	Windows Server 2008R2 SP1
SQL Server Edition	SQL Server 2012 SP1 Enterprise

(continued)

Configuration	Value
SQL Server Features	Database Engine Services Management Tools - Complete
SQL Server Instance Name	MSSQLSERVER (Default instance)
Computer name	EVDL-SQL-AG02
Role	Mirror (mirroring) Secondary (AlwaysOn) Failover Cluster node
vCPUs	2
Architecture	64-bit
Memory	2 GB
Storage	25 GB System Drive C:\ (SSD) 20 GB Data Drive D:\ (SSD)
Data Drive Layout	D:\Data, MDF Files D:\Log, LDF Files D:\Backup, Backup files
Operating System	Windows Server 2008R2 SP1
SQL Server Edition	SQL Server 2012 SP1 Enterprise
SQL Server Features	Database Engine Services Management Tools - Complete
SQL Server Instance Name	MSSQLSERVER (Default instance)

In-Memory OLTP Test Machine

This is the configuration of the virtual machine I used during Chapter 12, “In-Memory OLTP-Related Wait Types.”

Configuration	Value
Computer name	EVDL-SQL2014-01
vCPUs	2
Architecture	64-bit
Memory	4 GB
Storage	20 GB System Drive C:\ (SSD) 20 GB Data Drive D:\ (SSD)

(continued)

Configuration	Value
Data Drive Layout	D:\Data, MDF Files D:\Log, LDF Files D:\Backup, Backup files
Operating System	Windows Server 2008R2 SP1
SQL Server Edition	SQL Server 2014 Enterprise
SQL Server Features	Database Engine Services Management Tools - Complete
SQL Server Instance Name	MSSQLSERVER (Default instance)

APPENDIX II



Spinlocks

Spinlocks are described by Microsoft as “lightweight synchronization primitives.” The description looks a lot like the one used for latches, which are described as “lightweight synchronization objects.” This is no coincidence, as spinlocks and latches have a lot in common and both are used to serialize access to internal data structures. Both latches and spinlocks are used when access to objects needs to be held for a very short amount of time.

While spinlocks and latches have an identical purpose, there is one large difference between them. Whenever you cannot acquire a latch because there is another, incompatible latch already in place, for example, your request is forced to wait, and it will leave the processor and get returned to the Waiter List (the request receives the “SUSPENDED” state). It is then forced to wait inside the Waiter List until the latch can get acquired, and then it moves through the Runnable queue until it can finally get back on the processor. Because latches are treated like a resource for query execution, they are closely related to Wait Statistics. SQL Server even records the time it has been waiting on acquiring different latch types and classes, which we discussed in Chapter 8, “Latch-Related Wait Types.” There is a relatively large overhead associated with latches, because if a latch cannot be obtained immediately, it has to move through the different phases of the scheduler again before the request can acquire its latch and get executed on the processor.

Spinlocks work very differently than latches, because whenever a spinlock has to wait out another spinlock already in place before it can get placed itself, the thread does not have to leave the processor. Instead, a spinlock will “spin” until it can be acquired. Figure AII-01 shows the difference between latches and spinlocks whenever one has to wait before it can get acquired.

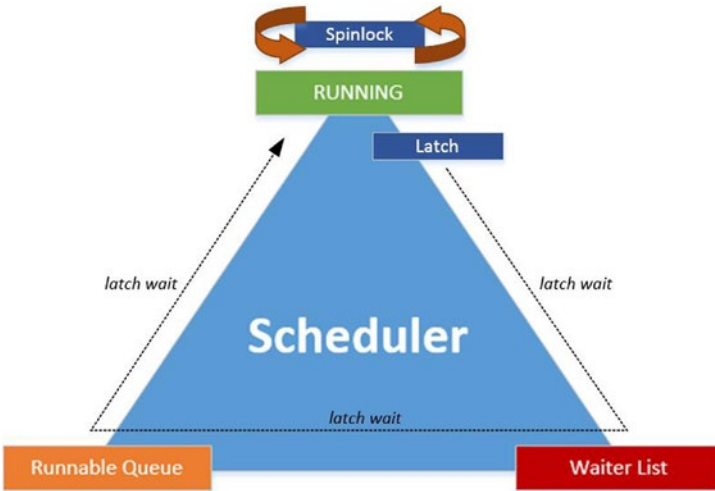


Figure AII-01. Spinlocks and latches and wait phases

The main advantage of using spinlocks instead of latches to synchronize thread access is that spinlocks are even “lighter” synchronization objects than latches. Latches cause extra context switching to occur whenever a latch has to wait before it can get acquired. Spinlocks do not cause context switching because they will never move away from the processor. Because spinlocks do not cause context switching, they are used to protect those areas of SQL Server that are used most intensely. Spinlocks are not the Holy Grail in protecting access to data structures, however. Because they never move away from the processor, they consume processor time, even when they are waiting. To avoid spinlocks consuming too much processor time, every *x* time around the spinlock will stop spinning and sleep. The interval of the spinlock sleep is calculated by an internal algorithm.

On very busy systems, where many spinlocks are used, it is possible to encounter a phenomena called spinlock contention. If the spinlock contention gets bad enough, you can notice an increase in processor time that can be difficult to troubleshoot, since this will not always show by analyzing Wait Statistics.

Thankfully, just like latches, there is a spinlock DMV inside SQL Server that tracks the specific spinlock classes (262 in SQL Server 2014), the amount of time a spinlock had to wait before it could get acquired, and the total number of spins that occurred for that spinlock class. We can access this information by querying the `sys.dm_os_spinlock_stats` DMV like the query here:

```
SELECT *
FROM sys.dm_os_spinlock_stats
ORDER BY spins DESC;
```

This returns results like those shown in Figure [AII-02](#).

	name	collisions	spins	spins_per_collision	sleep_time	backoffs
1	LOCK_HASH	29	8250	284.4828	0	2
2	RESQUEUE	6	2000	333.3333	0	1
3	SOS_CACHSTORE	7	1750	250	0	0
4	BACKUP_CTX	5	1250	250	0	0
5	SOS_SUSPEND_QUEUE	2	500	250	0	0
6	BLOCKER_ENUM	1	250	250	0	0
7	SOS_TLIST	1	250	250	0	0
8	SOS_RW	1	250	250	0	0
9	SOS_WAITABLE_ADDRESS_HASHBUCKET	0	0	0	0	0
10	SOS_ABORT_TASK	0	0	0	0	0

Figure AII-02. *sys.dm_os_spinlock_stats*

The columns returned by the `sys.dm_os_spinlock_stats` are described in the list here:

- `name`. Shows the name of the spinlock class.
- `collisions`. Returns the amount of time this spinlock class encountered a wait event because another spinlock was already in place.
- `spins`. When a spinlock has to wait it performs a spin. The spins column shows the amount of times spins occurred for this specific spinlock class. You can think of a spin as the amount of time the spinlock had to wait before it could get acquired.
- `spins_per_collision`. The average amount of spins per collision.
- `sleep_time`. Time that was spent sleeping for this spinlock class.
- `backoffs`. The number of times a spinlock went to sleep to allow other threads to use the processor.

While all the columns returned by the `sys.dm_os_spinlock_stats` DMV provide valuable information, the `backoffs` column can be the most interesting when you are suspecting a case of spinlock contention. If you notice very high CPU usage and cannot directly correlate the high CPU usage with queries or specific Wait Types, but the amount of backoffs for a specific spinlock class is very high and increasing quickly, you might have a case of spinlock contention occurring.

Spinlock contention is difficult to troubleshoot since it can have a very large number of causes. Also, information about specific spinlock classes is often lacking, increasing the difficulty of troubleshooting spinlock contention. One method that you can use during the analysis of spinlock contention is building a baseline of the `sys.dm_os_spinlock_stats` DMV by capturing the contents of the DMV at a specific interval, like I described in Chapter 3, “Building a Solid Baseline.” This baseline can give you valuable insight into the usage of spinlocks inside your SQL Server instance. Another great tool to diagnose spinlock contention is Extended Events. By using Extended Events you can trace various spinlock-related events, like spinlock backoffs.

To truly analyze why spinlock-class contention is occurring, you will have to dive even deeper by debugging SQL Server memory dumps and looking through the call stack to find what spinlock class is being accessed. Debugging SQL Server memory dumps to identify spinlock contention is beyond the scope of this book and requires a deep knowledge of the inner workings of SQL Server. Thankfully, there is a free Microsoft whitepaper available on spinlock contention that can give you a few pointers for what to do when dealing with spinlock contention. You can get the whitepaper at the following URL: <https://www.microsoft.com/en-us/download/details.aspx?id=26666>.

APPENDIX III



Latch Classes

Latch class	Books Online Description	Additional Information
ALLOC_CREATE_RINGBUF	Used internally by SQL Server to initialize the synchronization of the creation of an allocation ring buffer.	Used when creating a ring buffer. A ring buffer briefly holds internal event information in memory and is used for diagnostics.
ALLOC_CREATE_FREESPACE_CACHE	Used to initialize the synchronization of internal freespace caches for heaps.	Allocates free space for heaps (tables without a clustered index)
ALLOC_CACHE_MANAGER	Used to synchronize internal coherency tests.	
ALLOC_FREESPACE_CACHE	Used to synchronize access to a cache of pages with available space for heaps and binary large objects (BLOBs). Contention on latches of this class can occur when multiple connections try to insert rows into a heap or BLOB at the same time. You can reduce this contention by partitioning the object. Each partition has its own latch. Partitioning will distribute the inserts across multiple latches.	

(continued)

Latch class	Books Online Description	Additional Information
ALLOC_EXTENT_CACHE	Used to synchronize the access to a cache of extents that contains pages that are not allocated. Contention on latches of this class can occur when multiple connections try to allocate data pages in the same allocation unit at the same time. This contention can be reduced by partitioning the object of which this allocation unit is a part.	
ACCESS_METHODS_DATASET_PARENT	Used to synchronize child dataset access to the parent dataset during parallel operations.	Used together with the ACCESS_METHODS_SCAN_RANGE_GENERATOR latch class during parallel operations to distribute the work among multiple threads.
ACCESS_METHODS_HOBT_FACTORY	Used to synchronize access to an internal hash table.	
ACCESS_METHODS_HOBT	Used to synchronize access to the in-memory representation of a HoBt.	
ACCESS_METHODS_HOBT_COUNT	Used to synchronize access to a HoBt page and row counters.	Used for page and row count deltas for heaps and B-trees.
ACCESS_METHODS_HOBT_VIRTUAL_ROOT	Used to synchronize access to the root page abstraction of an internal B-tree.	Used when accessing metadata regarding the index's root page. See Chapter 8, "Latch-Related Wait Types," for an example.
ACCESS_METHODS_CACHE_ONLY_HOBT_ALLOC	Used to synchronize worktable access.	Used for synchronizing access to transparent, temporary tables that are created during query execution.
ACCESS_METHODS_BULK_ALLOC	Used to synchronize access within bulk allocators.	
ACCESS_METHODS_SCAN_RANGE_GENERATOR	Used to synchronize access to a range generator during parallel scans.	

(continued)

Latch class	Books Online Description	Additional Information
ACCESS_METHODS_KEY_RANGE_GENERATOR	Used to synchronize access to read-ahead operations during key-range parallel scans.	
APPEND_ONLY_STORAGE_INSERT_POINT	Used to synchronize inserts in fast append-only storage units.	
APPEND_ONLY_STORAGE_FIRST_ALLOC	Used to synchronize the first allocation for an append-only storage unit.	
APPEND_ONLY_STORAGE_UNIT_MANAGER	Used for internal data structure access synchronization within the fast append-only storage unit manager.	
APPEND_ONLY_STORAGE_MANAGER	Used to synchronize shrink operations in the fast append-only storage unit manager.	
BACKUP_RESULT_SET	Used to synchronize parallel backup result sets.	
BACKUP_TAPE_POOL	Used to synchronize backup tape pools.	
BACKUP_LOG_REDO	Used to synchronize backup log redo operations.	
BACKUP_INSTANCE_ID	Used to synchronize the generation of instance IDs for backup performance monitor counters.	
BACKUP_MANAGER	Used to synchronize the internal backup manager.	
BACKUP_MANAGER_DIFFERENTIAL	Used to synchronize differential backup operations with DBCC.	
BACKUP_OPERATION	Used for internal data structure synchronization within a backup operation, such as database, log, or file backup.	
BACKUP_FILE_HANDLE	Used to synchronize file open operations during a restore operation.	

(continued)

Latch class	Books Online Description	Additional Information
BUFFER	<p>Used to synchronize short-term access to database pages. A buffer latch is required before reading or modifying any database page. Buffer latch contention can indicate several issues, including hot pages and slow I/Os.</p> <p>This latch class covers all possible uses of page latches. <code>sys.dm_os_wait_stats</code> makes a difference between page latch waits that are caused by I/O operations and read and write operations on the page.</p>	<p>Directly related to buffer latches. When seeing higher-than-expected Wait Times, check if you are running into buffer latch-related contention.</p>
BUFFER_POOL_GROW	<p>Used for internal buffer manager synchronization during buffer pool grow operations.</p>	
DATABASE_CHECKPOINT	<p>Used to serialize checkpoints within a database.</p>	
CLR_PROCEDURE_HASHTABLE	<p>Internal use only.</p>	
CLR_UDX_STORE	<p>Internal use only.</p>	
CLR_DATAT_ACCESS	<p>Internal use only.</p>	
CLR_XVAR_PROXY_LIST	<p>Internal use only.</p>	
DBCC_CHECK_AGGREGATE	<p>Internal use only.</p>	
DBCC_CHECK_RESULTSET	<p>Internal use only.</p>	
DBCC_CHECK_TABLE	<p>Internal use only.</p>	
DBCC_CHECK_TABLE_INIT	<p>Internal use only.</p>	
DBCC_CHECK_TRACE_LIST	<p>Internal use only.</p>	
DBCC_FILE_CHECK_OBJECT	<p>Internal use only.</p>	
DBCC_PERF	<p>Used to synchronize internal performance monitor counters.</p>	
DBCC_PFS_STATUS	<p>Internal use only.</p>	
DBCC_OBJECT_METADATA	<p>Internal use only.</p>	
DBCC_HASH_DLL	<p>Internal use only.</p>	
EVENTING_CACHE	<p>Internal use only.</p>	

(continued)

Latch class	Books Online Description	Additional Information
FCB	Used to synchronize access to the file control block.	
FCB_REPLICA	Internal use only.	
FGCB_ALLOC	Use to synchronize access to round-robin allocation information within a filegroup.	
FGCB_ADD_REMOVE	Use to synchronize access to filegroups for ADD and DROP file operations.	Latch is used when adding or removing files inside a filegroup, or when a file grows. Check auto-growth configuration if you are running into contention.
FILEGROUP_MANAGER	Internal use only.	
FILE_MANAGER	Internal use only.	
FILESTREAM_FCB	Internal use only.	
FILESTREAM_FILE_MANAGER	Internal use only.	
FILESTREAM_GHOST_FILES	Internal use only.	
FILESTREAM_DFS_ROOT	Internal use only.	
LOG_MANAGER	Internal use only.	Indicates transaction-log growth because the log could not be cleared or truncated.
FULLTEXT_DOCUMENT_ID	Internal use only.	
FULLTEXT_DOCUMENT_ID_TRANSACTION	Internal use only.	
FULLTEXT_DOCUMENT_ID_NOTIFY	Internal use only.	
FULLTEXT_LOGS	Internal use only.	
FULLTEXT_CRAWL_LOG	Internal use only.	
FULLTEXT_ADMIN	Internal use only.	
FULLTEXT_AMDIN_COMMAND_CACHE	Internal use only.	
FULLTEXT_LANGUAGE_TABLE	Internal use only.	
FULLTEXT_CRAWL_DM_LIST	Internal use only.	
FULLTEXT_CRAWL_CATALOG	Internal use only.	
FULLTEXT_FILE_MANAGER	Internal use only.	

(continued)

Latch class	Books Online Description	Additional Information
DATABASE_MIRRORING_REDO	Internal use only.	
DATABASE_MIRRORING_SERVER	Internal use only.	
DATABASE_MIRRORING_CONNECTION	Internal use only.	Responsible for controlling the message flow between database mirrors.
DATABASE_MIRRORING_STREAM	Internal use only.	
QUERY_OPTIMIZER_VD_MANAGER	Internal use only.	
QUERY_OPTIMIZER_ID_MANAGER	Internal use only.	
QUERY_OPTIMIZER_VIEW_REP	Internal use only.	
RECOVERY_BAD_PAGE_TABLE	Internal use only.	
RECOVERY_MANAGER	Internal use only.	
SECURITY_OPERATION_RULE_TABLE	Internal use only.	
SECURITY_OBJPERM_CACHE	Internal use only.	
SECURITY_CRYPTO	Internal use only.	
SECURITY_KEY_RING	Internal use only.	
SECURITY_KEY_LIST	Internal use only.	
SERVICE_BROKER_CONNECTION_RECEIVE	Internal use only.	
SERVICE_BROKER_TRANSMISSION	Internal use only.	
SERVICE_BROKER_TRANSMISSION_UPDATE	Internal use only.	
SERVICE_BROKER_TRANSMISSION_STATE	Internal use only.	
SERVICE_BROKER_TRANSMISSION_ERRORS	Internal use only.	
SSBXmitWork	Internal use only.	
SERVICE_BROKER_MESSAGE_TRANSMISSION	Internal use only.	
SERVICE_BROKER_MAP_MANAGER	Internal use only.	

(continued)

Latch class	Books Online Description	Additional Information
SERVICE_BROKER_HOST_NAME	Internal use only.	
SERVICE_BROKER_READ_CACHE	Internal use only.	
SERVICE_BROKER_WAITFOR_MANAGER	Internal use only.	
SERVICE_BROKER_WAITFOR_TRANSACTION_DATA	Internal use only.	
SERVICE_BROKER_TRANSMISSION_TRANSACTION_DATA	Internal use only.	
SERVICE_BROKER_TRANSPORT	Internal use only.	
SERVICE_BROKER_MIRROR_ROUTE	Internal use only.	
TRACE_ID	Internal use only.	
TRACE_AUDIT_ID	Internal use only.	
TRACE	Internal use only.	
TRACE_CONTROLLER	Internal use only.	Related to SQL Trace. More information about SQL Trace can be found at https://msdn.microsoft.com/en-us/hh245121.aspx . Seeing contention on this latch class can mean too many traces are running at the time.
TRACE_EVENT_QUEUE	Internal use only.	
TRANSACTION_DISTRIBUTED_MARK	Internal use only.	
TRANSACTION_OUTCOME	Internal use only.	
NESTING_TRANSACTION_READONLY	Internal use only.	
NESTING_TRANSACTION_FULL	Internal use only.	
MSQL_TRANSACTION_MANAGER	Internal use only.	
DATABASE_AUTONAME_MANAGER	Internal use only.	

(continued)

Latch class	Books Online Description	Additional Information
UTILITY_DYNAMIC_VECTOR	Internal use only.	
UTILITY_SPARSE_BITMAP	Internal use only.	
UTILITY_DATABASE_DROP	Internal use only.	
UTILITY_DYNAMIC_MANAGER_VIEW	Internal use only.	
UTILITY_DEBUG_FILESTREAM	Internal use only.	
UTILITY_LOCK_INFORMATION	Internal use only.	
VERSIONING_TRANSACTION	Internal use only.	
VERSIONING_TRANSACTION_LIST	Internal use only.	
VERSIONING_TRANSACTION_CHAIN	Internal use only.	
VERSIONING_STATE	Internal use only.	
VERSIONING_STATE_CHANGE	Internal use only.	
KTM_VIRTUAL_CLOCK	Internal use only.	

Index

■ A

- ASYNC_IO_COMPLETION Wait Type
 - backup-related waits, 99–100
 - Books Online (BOL), 98
 - definition, 98
 - instant file initialization, 101
 - lowering, 99–102
 - perform volume
 - maintenance tasks, 100
 - storage-related Wait Types, 102–103
 - storage subsystem, 97
 - sys.dm_os_wait_stats DMV, 98–99
 - visual representation, 98
- ASYNC_NETWORK_IO Wait Type
 - application requests, 106
 - graphical representation, 103
 - lowering, 104–106
 - modification, 105
 - sys.dm_os_wait_stats DMV, 104
 - task manager network utilization, 105
- Atomicity characteristics, 145–146

■ B

- Background processes
 - benign, 247
 - CHECKPOINT_QUEUE, 248–251
 - DIRTY_PAGE_POLL, 251–253
 - LAZYWRITER_SLEEP, 253–255
 - miscellaneous, 247
 - MSQL_XP, 255–257
 - OLEDB Client Provider, 257–259
 - TRACEWRITE, 259–266
 - WAITFOR, 266–267
- BACKUPBUFFER Wait Type
 - backup/restore operation, 136, 140
 - BufferCount parameter, 138
 - database backup, 137
 - generating BACKUPBUFFER waits, 138
 - lowering, 138–139
 - reader and writer
 - backup buffers, 136–137
- BACKUPIO Wait Type
 - backup/restore operation, 140, 142
 - internals, 140–141
 - lowering, 142
 - modified backup query, 141
 - sys.dm_os_wait_stats DMV, 141–142
- Backups-related Wait Types
 - BACKUPBUFFER, 136–140
 - BACKUPIO, 140–142
 - BACKUPTHREAD, 142–144
 - RTO and RPO, 135
 - SQL Server
 - backup strategy, 135
- BACKUPTHREAD Wait Type
 - AdventureWorks
 - database, 143–144
 - backup/restore operation, 142
 - lowering, 144
 - relation, 143
 - summary, 144
- Baseline analysis
 - comparison graph, 64
 - disk-read latency, 64–65
 - flowchart, 61
 - measurements, 66–67
 - metrics, 60
 - PAGEIOLATCH_SH, 62–63
 - performance-analysis
 - flowchart, 65–66
 - sys.dm_os_waiting_tasks DMV, 62
- Buffer latches, 174

■ **C**

- Checkpoint file pairs (CFPs)
 - abstract view, 271
 - data file and row data, 271
 - memory-optimized tables, 270–271
 - merge operation, 272–273
 - multiple data files, 271–272
 - offline checkpoint thread, 272
 - transaction-commit
 - timestamps, 271–272
 - WAIT_XTP_HOST_WAIT, 278
- CHECKPOINT_QUEUE Wait Type
 - operations, 251
 - SQL Server instances, 248
 - sys.dm_os_wait_stats and sys.dm_os_waiting_tasks DMV, 249–250
 - sys.dm_os_wait_stats DMV, 250
- CMEMTHREAD Wait Type
 - lowering, 107–108
 - memory objects, 106, 108
 - Mutex object, 106–107
 - mutual exclusions, critical sections
 - and thread safety, 106
 - query procedure cache, 108
 - SQL Server mini-dump, 107
- Common Language Runtime (CLR), 243
- Consistency characteristics, 145–146
- CPU-related Wait Types
 - CXPACKET, 71–78
 - processors, 71
 - SOS_SCHEDULER_YIELD, 78–85
 - THREADPOOL, 85–96
- CXPACKET Wait Type
 - configuration of parallelism, 72–73
 - CPU-related, 71
 - execution plan, 75
 - lowering parallelism
 - configuration, 74–76
 - parallel queries, 71–72
 - query execution, 78
 - SELECT operation, 75–76
 - skewed workload, 77–78
 - threads, 72

■ **D**

- Database administrator (DBA), 29
- Database Management Systems (DBMSs), 5
- DBMIRROR_SEND Wait Type
 - AdventureWorks database, 211
 - asynchronous mode, 210

- communications layer, 211
- database mirroring
 - monitor, 208, 213–214
 - lowering, 213–214
 - Mirror_Test table, 211–212
 - modification transaction, 209–210
 - results, 212
 - synchronous mirror mode, 208–209
- Dedicated Administrator
 - Connection (DAC), 96
- Detect Waits
 - blocking information, 30
 - query results, 30
 - scenario, 29
 - sys.dm_exec_sessions, 30–31
 - sys.dm_os_waiting_tasks, 29
 - Wait Statistics flowchart, 32
- DIRTY_PAGE_POLL Wait Type
 - checkpoint feature, 251
 - configuration option, 252
 - dirty pages, 253
 - indirect checkpoints, 251–252
 - SQL Server 2012, 251
 - summary, 253
- Disk-based tables
 - isolation, 273–274
 - SQL Server instance, 270
- Durable characteristics, 145–146
- Dynamic Management Views (DMVs), 4
 - detect Waits, 29–32
 - extended events, 34–45
 - perfmon command, 33–34
 - sys.dm_exec_requests, 26–29
 - sys.dm_os_waiting_tasks, 23–26
 - sys.dm_os_wait_stats, 21–23

■ **E, F, G**

- Extended events
 - ALTER EVENT SESSION
 - command, 43
 - capture wait statistics
 - information, 35–45
 - configuration, 36–37
 - execution, 42–43
 - filter, 39–40
 - live Wait Statistics information, 41–42
 - monitor-event selection, 37–38
 - Person.Person table, 41
 - query results, 44–45
 - returns files as rows, 43–44
 - sessions folder, 35–36

SQL Server Profiler, 34–35
 sql_text global field, 38–39
 start extended event session, 40–41
 sys.dm_xe_map_values, 35
 types, 35

■ **H**

HADR_LOGCAPTURE_WAIT and HADR_WORK_QUEUE Wait Types
 AlwaysOn Availability Groups, 214
 asynchronous-commit mode, 215
 availability modes, 215
 database mirroring
 configuration, 214–215
 primary and secondary
 replicas, 215–216
 summary, 216
 synchronous-commit mode, 215
 HADR_SYNC_COMMIT Wait Type
 AdventureWorks database, 218
 AlwaysOn operation mode, 220
 AO_Test table, 219
 column bar, 220–221
 dashboard option, 220
 generation inside, 217
 lowering, 220–222
 perfmon counters, 221–222
 primary and secondary replicas, 219
 secondary replica, 222
 synchronous replication
 mode, 217, 223
 sys.dm_os_wait_stats DMV, 218
 High-availability and
 disaster-recovery (HA/DR)
 configurations, 207
 DBMIRROR_SEND, 208–214
 HADR_LOGCAPTURE_WAIT and
 HADR_WORK_QUEUE, 214–216
 HADR_SYNC_COMMIT, 217–223
 REDO_THREAD_
 PENDING_WORK, 223–225
 virtual machines, 207

■ **I, J, K**

INDEXPROPERTY function, 195
 In-Memory OLTP-related Wait Types
 checkpoint file pairs, 270–273
 disk-based tables, 270
 feature analysis, 269

isolation level, 273–274
 memory-optimized tables, 270
 traditional tables, 270
 transaction log changes, 274
 WAIT_XTP_CKPT_CLOSE, 279–281
 WAIT_XTP_HOST_WAIT, 274–278
 WAIT_XTP_OFFLINE_CKPT_
 NEW_LOG, 281–283

Internet Information Services (IIS), 229

IO_COMPLETION Wait Type
 database-related actions, 109
 lowering, 111
 storage-related actions, 109, 111
 transaction log backup, 109–110

IO latches, 174

IO-related Wait Types
 ASYNC_IO_COMPLETION, 97–103
 ASYNC_NETWORK_IO, 103–106
 CMEMTHREAD, 106–109
 IO_COMPLETION, 109–111
 LOGBUFFER and
 WRITELOG, 111–116
 RESOURCE_SEMAPHORE, 116–122
 RESOURCE_SEMAPHORE_
 QUERY_COMPILE, 122–128
 SLEEP_BPOOL_FLUSH, 128–133
 WRITE_COMPLETION, 133–134

Isolation levels

 characteristics, 145–146
 definition, 149
 dirty reads, 150
 inconsistent analysis, 150
 locking behavior, 150–151
 phantom reads, 150
 read uncommitted values, 149–150
 serializable, 150
 transaction, 151–152

■ **L**

Latches

 classes, 293–300
 spinlocks, 289–290

Latch-related Wait Types

 compatibility matrix, 173
 description, 171
 latch definition, 172
 LATCH_[xx], 189–196
 locking and blocking, 171
 modes, 172–173
 PAGEIOLATCH_SH occur, 173–174

Latch-related Wait Types (*cont.*)

- PAGEIOLATCH_[xx], 197–205
- page-latch contention, 176–179
- PAGELATCH_[xx], 180–189
- sys.dm_os_wait_stats, 174–176
- transactions, 171–172
- wait occurs, 173–175

LATCH_[xx] Wait Type

- B-tree index structure, 192–193
- buffer cache, 189
- INDEXPROPERTY function, 195
- lowering, 196
- memory area, 190
- navigation of B-tree, 193–194
- non-clustered index, 191
- resource_description
 - column, 191–192
- results, 192
- sys.dm_db_index_physical_stats
 - results, 195
- sys.dm_os_wait_stats and sys.dm_os_latch_stats, 190
- test contention table, 191
- TRUNCATE command, 194–195

LAZYWRITER_SLEEP Wait Type

- checkpoint and
 - lazywriter processes, 254
- flowchart, 254
- lazywriter process shares, 253–254
- SQL Server process, 255

LCK_M_I[xx] Wait Type

- COMMIT command, 165
- Intent lock, 163–164, 167
- locking hierarchy, 164–165
- lowering, 166
- SELECT statement, 165
- Shared Intent Update, 164
- Shared with Intent Exclusive, 164
- sys.dm_os_waiting_tasks DMV, 165
- Update Intent Exclusive, 164

LCK_M_SCH_S &

LCK_M_SCH_M Wait Types

- COMMIT/ROLLBACK
 - command, 168
- lowering, 170
- Schema locks, 167, 170
- Sch-S and Sch-M, 167–168
- SELECT query, 169
- sys.dm_os_waiting_tasks DMV, 168
- table modification query, 169

LCK_M_S Wait Type

- data modification query, 156
- lock-related Wait Type, 155
- lowering, 157–158
- SELECT command, 157
- Shared lock, 156
- transactions, 158

LCK_M_U Wait Types

- concurrent transactions, 159–160
- graphical representation, 159
- lowering LCK_M_U waits, 161
- Ostress utility, 160–161
- Person.Address table, 160
- session ID 74, 161
- Update lock mode, 158, 161

LCK_M_X Wait Type

- COMMIT statement, 163
- exclusive lock, 162–163
- HumanResources.Employee table, 163
- lock-related Wait Types, 162
- lower LCK_M_X Wait Times, 163
- SELECT statement, 162
- shared locks, 162

Lock-related Wait Type

- ACID properties, 145
- characteristics, 145
- concurrent lock situation, 147–148
- hierarchy, 148–149
- isolation characteristics, 145–146
- LCK_M_I[xx], 163–167
- LCK_M_S, 155–158
- LCK_M_SCH_S &
 - LCK_M_SCH_M, 167–170
- LCK_M_U, 158–161
- LCK_M_X, 162–163
- level of isolation, 149–152
- lock modes and
 - compatibility, 146–148
 - querying lock information, 152–155
 - relational database, 145

LOGBUFFER and WRITELOG Wait Types

- approaches, 115
- flowchart, 111–112
- insert row, 114–115
- log buffer and storage
 - subsystem, 111–112
- Ostress utility, 115
- processes transactions, 116
- transaction movement, 112–113
- trans_demo database, 113–114

- Lowering TRACEWRITE Waits
 - export SQL server profiler trace, 264
 - reader_spid column, 263–264
 - server-side tracing, 266
 - sp_trace_setstatus, 265–266
 - sys.traces, 265
 - sys.traces catalog view, 263
 - trace definition, 264–265

■ **M**

- Memory-optimized tables
 - Checkpoint file pairs, 270–273
 - SQL Server instance, 270
 - WAIT_XTP_HOST_WAIT, 274–276
- Merge operation, 272–273
- MSQL_XP Wait Type
 - lowering MSQL_XP Waits, 257
 - MARS, 255
 - PREEMPTIVE_
 - GETPROCADDRESS, 255
 - results of, 256
 - stored procedures, 256–258
 - sys.dm_os_wait_stats DMV, 256
- Multiple Active Result
 - Sets (MARS), 255–258

■ **N**

- Non-buffer latches, 174

■ **O**

- Object Linking and Embedding Database (OLEDB). *See* OLEDB Wait Type
- Offline checkpoint thread, 272
- OLEDB Wait Type
 - CHECKDB command, 258
 - Client Provider, 257
 - DBCC command, 257
 - internal processes, 259
 - lowering OLEDB Waits, 258–259
 - query result, 258
 - sys.dm_os_wait_stats DMV, 258

■ **P**

- PAGEIOLATCH_[xx] Wait Types
 - buffer cache latches, 198
 - data-page movement, 198–199
 - disk operations, 197

- IO performance script, 202–203
- latch-related Wait Type, 197, 205
- lowering, 200–205
- modification, 203–204
- monitoring storage
 - performance, 201–202
- results, 200
- storage subsystem, 197–198
- sys.dm_os_wait_stats DMV, 199–200
- usage of buffer cache, 205
- Page-latch contention, 176–179
- PAGELATCH_[xx] Wait Type
 - AdventureWorks database, 182
 - clustered index, 181–182
 - DBCC IND command, 183–184
 - graphical representation, 180
 - ID_Hash column, 187
 - last-page insert
 - contention, 181, 185–186
 - leaf level, 181–182
 - lowering, 184–189
 - non-partitioned index, 187–188
 - Ostress command, 182–183
 - page-latch contention, 180
 - partitions, 188–189
 - query results, 183
- Perfmom command, 33–34
- PREEMPTIVE_OS_
 - AUTHENTICATIONOPS
 - Wait Type
 - connection, 240
 - lowering, 241–242
 - mixed-mode authentication, 239
 - query, 240–241
 - results, 241
 - SQL Server connects, 239
 - summary, 242
 - Windows authentication
 - functions, 239
 - PREEMPTIVE_OS_ENCRYPTMESSAGE
 - and PREEMPTIVE_OS_DECRYPTMESSAGE Wait Types
 - account selection, 231
 - certificate selection, 232–233
 - connection properties, 233–234
 - encrypted connection, 234–235
 - encrypting and
 - decrypting messages, 228
 - features view, 229
 - Internet Information Services, 229
 - lowering, 235

PREEMPTIVE_OS_ENCRYPTMESSAGE
 and PREEMPTIVE_OS_DECRYPTMESSAGE Wait
 Types (*cont.*)
 permissions dialog opens, 231–232
 queries, 235
 self-signed certificate, 231
 Server Certificates view, 230
 SQL Server instance, 228–229
 summary, 235
 Transparent Data Encryption, 229

PREEMPTIVE_OS_GETPROCADDRESS Wait Type
 Common Language Runtime, 243
 extended stored procedures, 243,
 245–246
 master database, 244
 overview, 243
 query Wait Statistics, 245
 summary, 246
 xp_cmdshell, 244
 xp_getnetname, 245

PREEMPTIVE_OS_WRITEFILEGATHER Wait Type
 AdventureWorks database
 data file, 237
 default database file
 configuration, 236
 description, 236
 instant file initialization, 238
 lowering PREEMPTIVE_OS_WRITEFILEGATHER
 Waits, 237–238
 storage subsystem, 238
 Windows operating system, 236

Preemptive Wait Type
 graphical representation, 227

PREEMPTIVE_OS_AUTHENTICATIONOPS, 239–242

PREEMPTIVE_OS_ENCRYPTMESSAGE and
 DECRYPTMESSAGE, 228–235

PREEMPTIVE_OS_GETPROCADDRESS, 243–246

PREEMPTIVE_OS_WRITEFILEGATHER, 236–238

SQL Server, 227–228

Windows functions, 227–228

WRITEFILEGATHER function, 228

■ Q

Querying lock information
 resource_description column, 153
 sp_WhoIsActive, 155
 sp_WhoIsActive default results, 154
 sql_text output, 154
 stored procedure, 154
 sys.dm_tran_locks DMV, 152

Querying SQL Server Wait statistics.
See Dynamic Management
 Views (DMVs)

■ R

REDO_THREAD_PENDING_WORK Wait Type
 characteristic HADR_prefix, 223
 data modifications, 223–224
 Redo Thread, 225
 replication modes, 225
 secondary replica inside, 223–224
 summary, 225
 synchronous AlwaysOn group, 224

Relational Database Management
 System (RDBMS), 145

RESOURCE_SEMAPHORE_QUERY_COMPILE Wait Type
 actual execution plan properties, 125
 compilation memory, 123
 compilation-memory
 resource semaphore, 122–123
 concurrent connections, 124–125

DBCC MEMORYSTATUS
 command, 124, 126

gateways, 123
 lowering, 127–128
 memory shortage, 122
 resource semaphores, 122, 128
 sys.dm_os_waiting_tasks DMV, 126
 threshold value, 124

RESOURCE_SEMAPHORE Wait Type
 additional memory, 116
 AdventureWorks database, 117
 lowering, 121
 MemoryGrantInfo
 properties, 117–118
 parallel query, 118
 perfmon counter, 121

query-memory request, 116, 122
 required memory, 116
 resource semaphore, 116–117
 SELECT command, 119–120
 sys.dm_exec_query_resource_
 semaphores, 119–120
 sys.dm_os_waiting_tasks DMV, 119

■ **S**

Schedulers, tasks and worker threads
 DMVs, 15
 requests, 10
 results, 15–16
 schedulers, 14–15
 sessions, 9
 SQL Server scheduling, 8
 tasks, 10–11
 worker threads, 11–13
 Schema Stability (Sch-S) and Schema
 Modification (Sch-M), 167–168
 Serializable isolation level, 150
 Shared Intent Update (SIU), 164
 Shared with Intent Exclusive (SIX), 164
 SLEEP_BPOOL_FLUSH Wait Type
 CHECKPOINT command, 132
 checkpoint process, 128, 130–133
 database data file, 128–129
 data modification process, 129
 results, 132
 sys.dm_os_wait_stats, 132
 trans_demo database, 131
 types, 129–130
 Solid baseline
 analysis, 60–67
 (see Wait Statistics analysis)
 benefits, 49
 comparing measurements, 52
 definition, 48
 graph, 49–50
 impact baseline, 48
 intervals, 52
 iterative process, 48–49
 key information, 52
 metrics, 52
 performance problem, 47–48
 pitfalls, 51
 types and statistics, 50–51
 visualization, 49–50

SOS_SCHEDULER_YIELD Wait Times
 AdventureWorks database, 80
 CPU queries, 84–85
 lowering, 81–85
 Ostress, 82–83
 phases and queues, 79
 RUNNING state, 79
 situations, 81
 SQL Server instance, 85
 SQL Server scheduling, 78
 sys.dm_os_wait_stats DMV, 81
 worker thread, 79–80

Spinlocks
 advantage, 290
 class contention, 292
 difference, 289–290
 internal algorithm, 290
 latches, 289
 lightweight synchronization
 primitives, 289
 sys.dm_os_spinlock_stats, 290–291
 troubleshoot, 291

SQL Operating System (SQLOS)
 architecture, 6
 DBMS, 5
 preemptive scheduling approach, 7
 quantum, 7
 simplified architecture, 6–7
 UMS, 7–8

SQL server machine configurations
 HA/DR, 286–287
 In-Memory OLTP, 287–288
 Ostress, 285
 test machine, 285–286
 Sys.dm_exec_requests DMV, 26–29
 Sys.dm_os_waiting_tasks DMV
 description, 24–25
 query, 25–26
 results, 23–24
 Sys.dm_os_wait_stats, 174–176
 Sys.dm_os_wait_stats DMV, 21–23

■ **T**

TempDB files, 178–179
 Test machine
 HA/DR, 286–287
 In-Memory OLTP, 287–288
 SQL server machine, 285–286

THREADPOOL Wait Type

- CPU usage, 95
- CXPACKET waits, 93–94, 96
- dedicated administrator connection, 92
- error execution, 90–91
- formulas, 86
- gaining access, 91–92
- Ostress user sessions, 95
- parallelism lowering, 92–94
- processors configuration, 86–87
- quertes, 87
- SQL Server scheduling, 85, 88–91
- sys.dm_os_waiting_tasks, 90
- sys.dm_os_waiting_tasks DMV, 95
- tasks and worker threads, 89
- test machine, 87–88
- thread pool starvation, 86
- user connections, 94–96

TRACEWRITE Wait Type. *See also*

- Lowering TRACEWRITE Waits
- default event selection, 261
- management tools, 260
- query execution, 262–263
- SQL:BatchCompleted event, 262
- SQL Profiler trace, 259–260
- trace button, 261

Traditional tables

- commit timestamps, 271
- log changes, 274
- SQL Server instance, 270

Transparent Data Encryption (TDE), 229

■ **U, V**

- Update Intent Exclusive (UIX), 164
- User Mode Scheduling (UMS), 7

■ **W, X, Y, Z**

WAITFOR Wait Type, 266–267

Wait Statistics analysis

- baseline creation, 52
- capture method, 56–58
- database, 53
- datetime data type, 54
- delta method, 55
- disadvantage, 55
- job schedule, 59–60
- reset capture method, 56

- schedule measurements, 58–60
- screenshot, 58–59
- sys.dm_os_wait_stats DMV, 54
- T-SQL script, 56

Wait Statistics internals

- abstract view, 16
- configuration option, 3
- database administrators/developers, 3
- DBCC command, 4
- execution time calculation, 17–18
- history, 3–5
- milliseconds forward
 - with scheduler, 19
- phases and queues, 19
- Runnable Queue, 17
- RUNNING phase, 17
- running requests with scheduler, 18
- schedulers, tasks and worker threads, 8–16
- SQLOS, 5–8
- SUSPENDED, 16
- Wait Times and worker thread flow, 17

WAIT_XTP_CKPT_CLOSE Wait Types

- automatic/manual checkpoint, 279
- call stack, 280
- checkpoint operations, 281
- results, 280
- sys.dm_os_wait_stats DMV, 279–280

WAIT_XTP_HOST_WAIT

- CFPs creation, 278
- checkpoint command, 276–277
- database data and log files, 275
- definition, 274–275
- extended events session, 277
- hkengine.dll file, 277
- In-Memory filegroup files, 277–278
- SQL Server instance, 276
- sys.dm_os_wait_stats and sys.dm_os_waiting_tasks DMV, 276

WAIT_XTP_OFFLINE_CKPT_NEW_LOG Wait Type, 281–283

Worker threads

- formulas, 11
- processors page, 12
- state column, 13
- sys.dm_os_threads, 11
- sys.dm_os_workers, 13

WRITE_COMPLETION

- Wait Type, 133–134