



# SQL Server 2022 Revealed

A Hybrid Data Platform Powered by  
Security, Performance, and Availability

—  
**Bob Ward**

*Foreword by Rohan Kumar*

**Apress®**

# SQL Server 2022 Revealed

A Hybrid Data Platform Powered  
by Security, Performance,  
and Availability

**Bob Ward**

*Foreword by Rohan Kumar*

Apress®

# ***SQL Server 2022 Revealed: A Hybrid Data Platform Powered by Security, Performance, and Availability***

Bob Ward

North Richland Hills, TX, USA

ISBN-13 (pbk): 978-1-4842-8893-1

ISBN-13 (electronic): 978-1-4842-8894-8

<https://doi.org/10.1007/978-1-4842-8894-8>

Copyright © 2022 by Bob Ward

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.

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, Apress Media LLC: Welmoed Spahr

Acquisitions Editor: Jonathan Gennick

Development Editor: Laura Berendson

Coordinating Editor: Jill Balzano

Cover image designed by Freepik ([www.freepik.com](http://www.freepik.com))

Distributed to the book trade worldwide by Springer Science+Business Media LLC, 1 New York Plaza, Suite 4600, New York, NY 10004. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail [orders-ny@springer-sbm.com](mailto:orders-ny@springer-sbm.com), or visit [www.springeronline.com](http://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 [booktranslations@springernature.com](mailto:booktranslations@springernature.com); for reprint, paperback, or audio rights, please e-mail [bookpermissions@springernature.com](mailto:bookpermissions@springernature.com).

Apress titles 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 Print and eBook Bulk Sales web page at <http://www.apress.com/bulk-sales>.

Any source code or other supplementary material referenced by the author in this book is available to readers on GitHub (<https://github.com/Apress>). For more detailed information, please visit <http://www.apress.com/source-code>.

Printed on acid-free paper

*I've said this before, but it is worth repeating because it matters. This book is dedicated to the #sqlfamily, the most dedicated and passionate and largest technical community I know of in the world. Without you we would not have a product.*

# Table of Contents

<b>About the Author .....</b>	<b>xiii</b>
<b>About the Technical Reviewer .....</b>	<b>xv</b>
<b>Acknowledgments .....</b>	<b>xvii</b>
<b>Foreword .....</b>	<b>xix</b>
<b>Introduction .....</b>	<b>xxi</b>
<b>Chapter 1: Project Dallas Becomes SQL Server 2022.....</b>	<b>1</b>
Project Dallas .....	1
Becoming SQL Server 2022 .....	3
Announcing SQL Server 2022 .....	7
Private to Public Preview .....	11
The Path to General Availability .....	14
Introducing SQL Server 2022 .....	14
Built on a Foundation.....	15
Wheel of Power .....	16
Cloud Connected .....	17
Built-In Query Intelligence.....	18
Industry-Proven Database Engine .....	18
Data Virtualization and Object Storage .....	19
Enhancing T-SQL for Developers .....	19
Getting Started with SQL Server 2022 .....	20
How to Get SQL Server 2022 .....	20
Installing SQL Server 2022 .....	20
Learn All the Features and Editions.....	21
Learn About Pricing and Licensing.....	21
Get Training on SQL Server 2022.....	21

TABLE OF CONTENTS

- Go Deeper with Our Blog Series ..... 21
- Download Book Code and Samples..... 21
- A Cloud-Connected, Intelligent, and Industry-Proven Data Platform ..... 22
- Chapter 2: Install and Upgrade ..... 25**
  - How to Install SQL Server 2022 ..... 25
    - Prerequisites ..... 26
    - What Is Different for SQL Server 2022?..... 26
    - Other Installation Methods ..... 31
  - Setting Up the Azure Extension for SQL Server..... 32
    - What You Should Know First..... 33
    - Providing Values for the Feature Setup ..... 36
    - Connect to Azure After Setup..... 40
    - Removing the Azure Extension for SQL Server..... 41
  - Deploying on Other Platforms ..... 41
  - Side-by-Side and Multi-instance Installations ..... 42
  - How to Upgrade to SQL Server 2022 ..... 42
  - The Importance of dbcompat..... 42
  - Configuration..... 43
  - Easy to Install and Upgrade ..... 43
- Chapter 3: Connect Your Database to the Cloud ..... 45**
  - The Hybrid SQL Server ..... 45
    - What Is *Hybrid*? ..... 46
    - SQL Server Hybrid Over the Years ..... 46
    - The SQL Server 2022 Hybrid Lineup..... 47
  - Managed Disaster Recovery with Azure SQL Managed Instance..... 49
    - Project Chimera and DAG ..... 50
    - The Link Feature for Azure SQL Managed Instance..... 52
    - How It Works ..... 52
    - Using the Link Feature for Offline Disaster Recovery..... 55
    - Keep in Mind These Details ..... 71
    - The Future for the Link Feature for Azure SQL Managed Instance..... 72

Azure Synapse Link for SQL Server .....	73
What Is Synapse Link for SQL Server? .....	74
How Does Synapse Link Work? .....	74
Try Out Synapse Link for SQL Server.....	77
More Details About Synapse Link.....	108
Synapse Link Could Change Analytics for You.....	113
Azure Active Directory (AAD) Authentication.....	113
How Does AAD Authentication Work? .....	114
Setting Up and Using AAD Authentication.....	115
Microsoft Purview Policy Management.....	124
How Do Purview Access Policies Work?.....	125
Using Purview Access Policies .....	128
Connecting SQL to the World .....	144
<b>Chapter 4: Built-In Query Intelligence .....</b>	<b>145</b>
Built-In Query Intelligence in SQL Server 2022.....	147
The New Query Store .....	149
On by Default.....	150
Query Store Hints .....	152
Query Store Support for Secondary Replicas .....	154
Store for IQP .....	158
IQP Nextgen Defaults .....	159
Approximate Percentile .....	159
Optimized Plan Forcing.....	160
IQP Nextgen with dbcompat 140+ .....	169
Memory Grant Percentiles .....	172
Memory Grant Feedback Persistence.....	172
The Intelligent Query Processor.....	183
<b>Chapter 5: Built-In Query Intelligence Gets Even Better .....</b>	<b>185</b>
Parameter-Sensitive Plan (PSP) Optimization.....	185
What Is a Parameter-Sensitive Plan? .....	186
How It All Started.....	188

TABLE OF CONTENTS

- How Does PSP Optimization Work? ..... 189
- Let's See PSP Optimization in Action..... 192
- What Other Details About PSP Optimization Should I Know? ..... 208
- PSP Optimization Is a Powerful Innovation..... 211
- Cardinality Estimation (CE) Model Feedback ..... 211
  - What Is the CE Model Problem? ..... 212
  - Try Out an Exercise..... 214
  - Limits and More Details About CE Feedback..... 224
- Degree of Parallelism (DOP) Feedback ..... 225
  - Why Did We Build DOP Feedback? ..... 226
  - What Is DOP Feedback? ..... 227
  - How Does DOP Feedback Work? ..... 227
  - Let's Try DOP Feedback..... 230
  - What Else Should I Know About DOP Feedback?..... 240
- An Engine That Works for You ..... 242
- Chapter 6: The Meat and Potatoes of SQL Server ..... 245**
  - The Core Engine Is the Meat and Potatoes of SQL Server ..... 245
  - Security..... 246
    - Ledger for SQL Server ..... 246
    - Encryption Enhancements..... 265
    - Security Permission Enhancements..... 269
  - Performance and Scalability..... 272
    - Columnstore and Batch Mode Improvements ..... 273
    - Scalability Improvements ..... 275
    - “Hands-Free” tempdb ..... 278
    - More Concurrency Improvements ..... 288
  - Availability..... 289
    - Contained Availability Groups ..... 289
    - Other AG Enhancements..... 302
    - Recovery Enhancements ..... 304



Backup/Restore Enhancements .....	306
Multi-write Replication .....	311
Other Engine “Stuff” .....	312
XML Compression.....	312
In-Memory OLTP Memory Management .....	313
Auto-drop Statistics.....	313
Resumable Add Table Constraints .....	313
New Wait Types .....	314
New Extended Events.....	314
An Industry-Proven Engine .....	314
<b>Chapter 7: Data Virtualization and Object Storage .....</b>	<b>317</b>
Data Virtualization in SQL Server 2022 .....	318
REST, Azure Storage, and S3.....	318
Project Gravity Becomes Polybase v3 .....	319
Polybase v3 File Formats .....	321
Using the New Polybase v3 .....	322
Try Out Polybase v3.....	324
What Else Do You Need to Know.....	344
Backup and Restore with S3 Compatible Object Storage .....	344
How Does It Work? .....	345
Let’s Look at an Example.....	346
Migration from AWS.....	348
SQL Server Is a Data Hub .....	350
<b>Chapter 8: New Application Scenarios with T-SQL .....</b>	<b>351</b>
JSON Functions.....	352
Prerequisites for Exercises.....	353
JSON_ARRAY and ISJSON .....	353
JSON_OBJECT and JSON_PATH_EXISTS.....	357
T-SQL Surface Area .....	358
Prerequisites .....	359
DATETRUNC .....	359

TABLE OF CONTENTS

- WINDOW Clause..... 361
- GREATEST and LEAST ..... 363
- IS [NOT] DISTINCT FROM ..... 366
- STRING\_SPLIT..... 369
- TRIM Function Extensions ..... 370
- Bit T-SQL Functions ..... 372
- Time Series ..... 379
  - Prerequisites ..... 380
  - DATE\_BUCKET ..... 380
  - GENERATE\_SERIES ..... 382
  - Gap Filling with FIRST\_VALUE and LAST\_VALUE ..... 385
- T-SQL Is Alive and Well..... 388
- Chapter 9: SQL Server 2022 on Linux, Containers, and Kubernetes ..... 389**
  - SQL Server 2022 on Linux..... 390
    - What’s New for SQL Server 2022 ..... 390
    - Deploying SQL Server 2022 on Linux ..... 392
    - How to Connect and Use SQL Server 2022 on Linux..... 394
    - Optimizing SQL Server 2022 on Linux ..... 396
    - HADR for SQL Server 2022 on Linux..... 397
  - SQL Server 2022 Containers ..... 399
    - Why Containers?..... 399
    - Using SQL Server 2022 Containers..... 401
    - Building a Customized Container Image..... 404
    - The Container Switch Method ..... 404
  - SQL Server 2022 on Kubernetes ..... 405
    - Why k8s..... 406
    - Deploying SQL Server on k8s ..... 406
    - Connecting and Using SQL Server on k8s ..... 410
    - High Availability with k8s ..... 410
    - Helm Charts..... 412
  - The World of Linux, Containers, and k8s..... 412

<b>Chapter 10: SQL Server 2022 on Azure Virtual Machines .....</b>	<b>413</b>
What Is SQL Server on Azure Virtual Machines? .....	413
Planning for Deployment.....	414
The SQL Server IaaS Agent Extension.....	418
What Is the SQL Server IaaS Agent Extension? .....	418
IaaS Agent Extension Modes .....	419
Installing the Agent Extension .....	419
Details of the Agent Extension.....	420
Deploying SQL Server on an Azure Virtual Machine.....	421
Prerequisites .....	421
Steps to Deploy .....	421
Other Deployment Methods .....	431
Exploring and Connecting to the SQL Server Virtual Machine .....	432
Exploring Your Azure Virtual Machine in the Portal.....	432
Connecting to the Azure Virtual Machine.....	436
Exploring a SQL Virtual Machine in the Portal .....	438
Connecting to SQL Server on an Azure Virtual Machine .....	439
Migrating to SQL Server on an Azure Virtual Machine .....	440
Optimizing Performance .....	441
Virtual Machine Size.....	441
Storage Performance .....	442
SQL Best Practices Assessment .....	443
High Availability.....	445
Built-In Fault Tolerance with Azure.....	445
Failover Cluster Instance (FCI).....	446
Always On Availability Groups .....	447
Disaster Recovery .....	449
Storage Fault Tolerance.....	449
Backup Database Options .....	449
Using Azure Backup.....	450

TABLE OF CONTENTS

- Monitoring..... 450
  - Azure Monitor ..... 451
  - Virtual Machine Metrics and Logs ..... 451
  - Azure Insights..... 452
- Azure Is the Best Cloud for SQL Server..... 452
- Chapter 11: SQL Edge to Cloud ..... 455**
  - Develop Once, Deploy Anywhere..... 455
  - Azure SQL Edge..... 457
    - When to Use Azure SQL Edge ..... 458
  - SQL Server ..... 458
    - When to Use SQL Server..... 459
    - Azure Arc–Enabled SQL Server..... 460
  - Linux, Containers, and Kubernetes ..... 460
    - When to Use SQL Server on Linux ..... 460
    - When to Use SQL Server Containers ..... 461
    - When to Use SQL Server on Kubernetes ..... 462
  - SQL Server on Azure Virtual Machines..... 463
    - When to Use SQL Server on Azure Virtual Machines ..... 464
  - Azure SQL Managed Instance ..... 465
    - When to Use Azure SQL Managed Instance..... 466
  - Azure SQL Database..... 467
    - When to Use Azure SQL Database ..... 469
  - Azure Arc–Enabled SQL Managed Instance ..... 470
    - When to Use Azure Arc–Enabled SQL Managed Instance..... 472
  - SQL Is Everywhere You Need It ..... 472
- Index..... 475**

# About the Author



**Bob Ward** is a principal architect for the Microsoft Azure Data team, which owns the development for all of SQL and Azure SQL. Bob has worked for Microsoft for 29 years on every version of SQL Server shipped from OS/2 1.1 to SQL Server 2022, including Azure SQL and Azure Arc. He is a well-known speaker on SQL Server and Azure SQL, often presenting talks on new releases, internals, and specialized topics at events such as the PASS Summit, SQLBits, Azure Data Conference, VSLive, Microsoft Build, Microsoft Inspire, Microsoft Ignite, and many other events. You can follow him

on Twitter at [@bobwardms](https://twitter.com/bobwardms) or LinkedIn at <https://linkedin.com/in/bobwardms>.

Bob is the author of Apress books *Pro SQL Server on Linux*, *SQL Server 2019 Revealed*, and *Azure SQL Revealed*.

# About the Technical Reviewer



**Erin Stellato** is a senior program manager on the SQL Experiences team, helping advance tools that customers use daily with Azure SQL. She is passionate about data and chocolate, not always in that order. She previously worked as a consultant and was a Data Platform MVP and has been an active member of the SQL Server community as both a volunteer and speaker. Her areas of interest within the engine include Query Store, Extended Events, statistics, and performance tuning, and she also enjoys helping accidental/involuntary DBAs figure out how SQL Server works.

# Acknowledgments

I want to first thank God for sending his one and only son, Jesus, for the forgiveness of my sins, the immersion of the Holy Spirit, and the promise of everlasting life. Without my faith, I am nothing.

I want to thank my wife, Ginger, so much for her love and dedication to me for all of our years together but especially during the writing of this book. You patiently allowed me to work crazy hours and situations even on some of our vacation time to bring this book to life. I also want to thank my sons Troy and Ryan, and Ryan's wife Blair, my daughter-in-law. You all give me hope for the future as I see you all grow and exhibit grace and truth in everything you do.

This is my fourth book with Apress, and I want to personally thank Jonathan Gennick from Apress. Jonathan gave me my first chance at authoring when many publication companies turned me down. Thank you, Jonathan, for always supporting me and letting me "write my way." I also want to thank Jill Balzano from Apress whom I've never met in person but who always is so kind and professional despite all the crazy deadlines we try to meet.

I asked Erin Stellato to be my technical reviewer because she is one of the deepest experts on SQL and just an incredible person. Turns out it was a blessing when she joined Microsoft during the writing of the book as we could discuss confidential information. Erin, thank you for your thoroughness and great attitude as we all put immense pressure on you to review so many chapters late in the cycle.

There were so many people at Microsoft who supported my work on this book. But, first, I want to thank Joe Sack, Pedro Lopes, and James Rowland-Jones who no longer work at Microsoft but were a huge part of helping me craft the story of SQL Server 2022, which you see in this book.

At Microsoft I want to thank my leaders who support all of my efforts including Rohan Kumar, Peter Carlin, Asad Khan, and Sanjay Mishra. The true heroes of this book are all the people at Microsoft in the engineering team who helped me with all my questions and the review of complex topics and for giving me some great quotes. This list is long, but I have to call out everyone who helped because each of them

## ACKNOWLEDGMENTS

helped in a way that made the entire book and the product worth the effort. Thank you, Ajay Jagannathan, David Pless, Kevin Farlee, Derek Wilson, Kate Smith, Kendal Van Dyke, Sasha Nosov, Travis Wright, Hanuma Kodavalla, Naveen Prakash, Joachim Hammer, Chuck Heinzelman, Mine Tokus, Hugo Queiroz, Tim Chen, Dani Ljepava, Mladen Andzic, Andreas Wolter, Srđan Bozovic, Vlad Rodriguez, Mirek Sztajno, Dylan Gray, Devin Rider, Robert Dorr, Peter Byrne, Jack Li, Mike Ray, Van To, Lee Woods, Sarah Kaufman, Conor Cunningham, Tejas Shah, Amit Khandelwal, Pam Lahoud, Purvi Shah, Aditya Badramraju, Taryn Pratt, Lani O'Brien, Nancy McGrory, Ebru Ersan, Umachandar Jayachandran (UC), Nicholas Simmons, Sharanya Bhat, Dilan Galindo Reyna, Ryan Stonecipher, Ravinder Vuppula, Thierry Fevrier, Aashna Bafna, Panagiotis Antonopoulos, Pieter Vanhove, Ankit Mahajan, Swaroop Moida, Fang Hou, Balmukund Lakhani, Parag Paul, Jakub Szymaszek, Yuqing Li, Perry Skountrianos, Pratim Dasgupta, Milos Vucic, Nikolas Ogg, David Liao, Chandrashekhara Kadiam, and Brian Carrig.

I also want to thank some of my colleagues on the engineering team who worked tirelessly to help me with workshops and events. Thank you, Marisa Mathews and Rie Merritt. You all work so hard behind the scenes that no one understands the effort you put in. And a special thank-you to my colleague Buck Woody. You are a dear friend and one of the most talented people I know. You always make traveling on the road a fun adventure.

Thank you also to the Microsoft Mechanics team for helping us highlight the amazing technology of SQL. Thank you, Jeremy Chapman and Celine Allee. Being on your show is always a privilege and fun experience.

I also want to thank members of our marketing team who were great partners during the launch and release of SQL Server 2022 including Matthew Burrows, Miwa Monji, Eric Hudson, Sonya Waitman, Guy Schoonmaker (thank you for your deck wizardry), and Emilija Dufresne.

Thank you, everyone, even for the smallest email response that helped make this book possible.



# Foreword

Today's modern business landscape is dynamic and fast moving – with data playing a fundamental role. Organizations need to quickly gain insights from data to inform customer experiences and empower their employees. Currently, many companies operate in a hybrid environment, with data in the cloud and on-premise. That's why we are so excited about the upcoming release of SQL Server 2022, a key part of our evolving data and intelligence strategy. At Microsoft, we have made significant investments in SQL, creating the most connected SQL database ever with this newly launched product – and it's incorporated into a Microsoft data ecosystem that gives organizations a seamless way to connect across their databases, analytics, and governance platforms.

I can think of no person more qualified to take you through a deep dive into SQL Server 2022 than our principal architect, and 29-year Microsoft veteran, Bob Ward. Bob's depth of knowledge is second only to his enthusiasm and passion for the SQL community.

In his informational book, Bob delivers a closer look at new solutions that enable you to do more with your data and guide you to take advantage of new capabilities. Bob will introduce Azure Synapse Link for SQL Server 2022, which gets you from data to insights faster and easier than ever.

Bob will cover database engine innovations including built-in query intelligence, Ledger for SQL Server and Contained Availability Groups, as well as Azure Synapse Link for SQL, which combines data from on-prem and cloud data. This enables near-real-time analytics over your most important operational data, all without the need to build and manage complex packages and pipelines – or the need to write a single line of code for ETL movement. You'll learn about SQL Managed Instance for managed disaster recovery and use the power of built-in Distributed Availability Groups (DAGs) to seed and link databases to Azure.

Additionally, Bob shares easy-to-consume information on Microsoft Purview, which you can use to create a policy for access rights and then publish across many SQL Servers, allowing you to have a central governance hub.

## FOREWORD

We hope you enjoy reading about, exploring, and deploying SQL Server 2022 as much as we enjoyed developing it. We want to extend our sincere thanks to our incredibly engaged SQL community. Your feedback over the years has helped us make a powerful platform that will reduce complexity, accelerate innovation, and deliver insights to inform business decisions. Thank you for giving us an opportunity to support your efforts and innovations!

Rohan Kumar  
Microsoft Corporate Vice President, Azure Data

# Introduction

Like *SQL Server 2019 Revealed*, the manuscript for this book saw plenty of mileage. I was back on the road again in late 2021 and in 2022, and many of the chapters were written or polished in places like Las Vegas (multiple times); Orlando; London (including the London Underground); Charleston, South Carolina (where my son Troy lives); Greenville, South Carolina; Dallas (where my son Ryan lives); Poolville, Texas (where my wife’s family lives); Austin; Genesee, Colorado (where Ginger and I have a small mountain retreat); Siesta Key, Florida (thank you, Tom and Janet Grubish); Redmond, Washington (multiple times); Chicago; and Atlanta. This includes a number of times writing pages in hotels, airplanes, trains, and even the front seat on a road trip while Ginger drove. But I was always most effective finishing off a chapter or making edits in the confines of my home in North Richland Hills, Texas.

This book is a complete knowledge transfer from me to you on everything I learned, absorbed, tested, and exhaustively poured my brain into about SQL Server 2022 from its inception as project Dallas right up to the General Availability release of the product. And it is my honor and pleasure to do that. All those late nights and weekends of study, testing, looking at source code, talking to program managers and developers, and thinking of scenarios based on my 29 years of experience are all in this book. I tried to bring both the high-level “what and why” with deep “how it works” so you can learn SQL Server 2022 from all angles. The number of pages alone in the book is a testament toward that goal. I’ve also tried to provide a deep set of examples in almost every chapter. I want you to not only read about why SQL Server 2022 is special but try it out for yourself to see it in action. I think you will also love in almost every chapter quotes from our engineering team on their perspective on why they believe SQL Server 2022 is a great release.

If you like knowing the history and “behind the scenes” of SQL Server 2022, start with [Chapter 1](#) where you will also get an insight into the overall story of SQL Server 2022. If you want to understand differences in installing SQL Server 2022 from previous releases including what has been removed and new Azure connectivity options, take a look at [Chapter 2](#).

## INTRODUCTION

All of the other chapters can be read in succession or on their own as separate topics. If you want an immersion into how SQL Server is connected to Azure, Chapter 3 is a complete resource. There are plenty of screenshots in this chapter to show you the entire Azure experience. Even if you don't have an Azure subscription, check out how we are connecting to Azure for disaster recovery (DR), analytics, and security.

Perhaps you want to dive right into the engine. Built-in query intelligence is so vast I needed two chapters to cover it all. Chapters 4 and 5 include very interesting exercises you can try yourself all on your own laptop with Evaluation or Developer Edition.

Chapter 6 is "all engine," and I believe you will like the 25+ major features we have introduced in SQL Server 2022 including Ledger for SQL Server, "hands-free" tempdb, and Contained Availability Groups.

If you have heard of S3 object storage but are not sure how to use it, you will like Chapter 7. I'll show you how to access parquet files and delta tables from SQL Server using our Polybase technology. Or perhaps you are looking for a new place to host your backups, so I'll show you how to back up and restore native SQL backups to and from any S3-compatible storage provider.

I think many believe we have forgotten about the Transact-SQL (T-SQL) language. You might be surprised to see the new T-SQL functions and language enhancements we have poured into SQL Server 2022 in Chapter 8.

There are not many new features in SQL Server 2022 specific to Linux, but I want to make sure you know the fundamentals. In Chapter 9 I'll show the basics of deploying and using SQL Server on Linux, containers, and Kubernetes.

I believe Azure is the best cloud for SQL Server, so Chapter 10 is a journey for you to learn how to deploy, optimize, and manage SQL Server 2022 on an Azure Virtual Machine.

And finally I wanted a way to close out the book with how SQL Server has become a force everywhere you need it. So Chapter 11 is how SQL Server exists "edge to cloud" in ways you have never dreamed with a story of consistency and flexibility.

If you have liked reading my books before or maybe this is for the first time, I hope you enjoy the book in all its aspects. I write in a very conversational style, which many have told me helps them "visualize me speaking to them" as they have seen me present in various events both in person and online. I also want you to have resources, so you will see plenty of online references throughout the book for you to dive in deeper.

Anytime you write a book like this, you want readers to have the latest information. So keep up to date with the latest examples and errata for the book at <https://aka.ms/sql2022bookexamples>. Or check out my free workshop materials at <https://aka.ms/sql2022workshop>.

I love to hear from readers, so you can either give the book a rating on Amazon or email me directly at [bobward@microsoft.com](mailto:bobward@microsoft.com). I always want to know how you feel about the book or if you want to learn more about SQL Server 2022.

Bob Ward  
North Richland Hills, TX  
September 2022

## CHAPTER 1

# Project Dallas Becomes SQL Server 2022

SQL Server 2022 was already in the making before SQL Server 2019 was released. In this chapter of the book, I will walk you through the history of SQL Server 2022, as it began with a project that turned into one of the most successful releases in SQL Server history.

I'll also give an introduction into why you want to consider using SQL Server 2022 as a data platform that is cloud connected, intelligent, and industry proven.

My goal is that you will read this chapter, get excited, and be hungry to delve into the following chapters of the book to learn all the details you need to know about SQL Server 2022.

## Project Dallas

In December of 2019, I found myself celebrating the release of SQL Server 2019 and catching up with all my colleagues on the SQL Server team in building 43 on the Microsoft campus in Redmond. It was a heady time for SQL Server and me. I had just traveled some 50,000 miles over the last calendar year getting the message out on SQL Server 2019. I had launched my second book, *SQL Server 2019 Revealed* (<https://aka.ms/sql2019book>), telling the entire story of our latest SQL Server release. I felt proud of the release and was able to tell all the great customer feedback and stories to many in the engineering team. On that trip I spoke to a large crowd of all the engineers who worked on the release in an internal meeting organized by the famous Slava Oks, Distinguished Engineer at Microsoft and Lead Development Manager for the SQL Server 2019 release.

I spent some of my time in Redmond catching up with my longtime friend from CSS, Robert Dorr, who now worked for Slava. I spent time at lunch telling Bob all the tales from my travels and my experiences launching the product and authoring the book. SQL Server 2019 was tagged project Seattle by Slava, Travis Wright, and Tobias Ternstrom.

I told Bob at lunch, “Hey, I think you and I should come up with the project name for the next version of SQL Server.” Bob started a year after me, so together we represented 55+ years of experience with SQL Server. That day Bob (as he often does) was sporting a Dallas Cowboys jersey. I proposed the project name should be called Dallas. After lunch I stopped by Slava’s desk and pitched the idea. He told me, “It makes sense. This new version should be a tribute to the two ‘Bobs’ who have worked so hard for SQL Server over the years.” Slava sent out an email the next day to our teams, and thus the project name Dallas was born. Sometimes that is how project names are created at Microsoft.

Even though we had just shipped SQL Server 2019, discussions had already begun about what type of new functionality could be part of project Dallas. For me, my focus would change drastically. On this same trip, I had meetings with my manager, Asad Khan, Vice President of Program Management (PM) across all of Azure SQL and SQL Server, which would lead me to start working on Azure SQL. But I still had my ear to the ground on project Dallas.

Already in late December of 2019, my colleague Amit Banerjee was putting together a proposal for the project. As we turned the corner to start 2020, little did I or anyone know all of us would have to adapt to a worldwide pandemic. In all honesty, the pandemic accelerated work for our teams in the cloud, which led to a slowdown in planning for project Dallas. The good news is that based on our engineering systems, much of the work we design and build in the cloud just naturally applies to SQL Server; Intelligent Query Processing (IQP) and SQL Ledger are great examples.

The focus on the cloud worked out very well for me personally, as I was already on the path to work on Azure SQL. I didn’t realize that the pandemic would speed up that work in the form of a partnership with Anna Hoffman for the Azure SQL Workshop (<https://aka.ms/azuresqlworkshop>), which then led to the popular Azure SQL Fundamentals learning path (<https://aka.ms/azuresqlfundamentals>) and a huge series of videos called Azure SQL for Beginners (<https://aka.ms/azuresql4beginners>). And since we had moved to all digital for our content, why not a third book? Thus was born *Azure SQL Revealed* (<https://aka.ms/azuresqlbook>), which launched in the fall of 2020.

But what about project Dallas?

Throughout calendar year 2020, project Dallas work had started under the leadership of Slava on the development side and Amit for program management. One example is degree of parallelism (DOP) feedback under a project called Gaia. Gaia had big ambitions and is an endeavor we are still working on, but the DOP feedback concept

was born out of it, and you will hear more about this in Chapter 5 of the book. Other capabilities like Parameter-Sensitive Plan (PSP) optimization, cross-platform backups, buffer pool scan, Query Store hints, multi-write replication, SQL Ledger, T-SQL language improvements, and more were all under way. We were actively working through improvements for project Dallas but had not really landed on the full plan for a release. Again, the cool thing about how we do things now is that many of these features were designed for Azure SQL, but we knew because of our engine compatibility they could be fit into Dallas at any time.

As we moved further along calendar year 2020, it became clear to me that project Dallas would not make it into an actual release until calendar year 2022 at the earliest. I had meetings with Slava, Amit, Asad, and Travis Wright late in 2020 to discuss the plan. All of us agreed, along with our Vice President of Azure Data, Rohan Kumar, that we needed a new major version of SQL Server sooner rather than later. But we also agreed that trying to do anything sooner than 2022 would not result in a great release. In late calendar year 2020, I would hear for the first time the words *SQL Server 2022*.

## Becoming SQL Server 2022

When this book is released, I will have been at Microsoft 29 years, all working on SQL. Throughout these years I've seen many changes, reorganizations, and transitions, more than I can even remember. But two changes happened early in calendar year 2021 that surprised me and gave me concern for project Dallas. Slava Oks left the SQL team to work in Azure (note: Slava would eventually leave Microsoft in the fall of 2021), and Amit Banerjee left Microsoft. In addition, Travis Wright pivoted from a role in the program management leadership team to work as a development manager for Azure Arc. Losing someone like Slava was pretty hard for me personally. I have worked with Slava in SQL for 20 years. All the journeys we had together (Slava is the one who taught me debugging in case anyone ever wanted to know that small fact; I thought I knew what debugging was until I met Slava) from SQL Server 2005 to SQL Server on Linux and of course SQL Server 2019. I also would deeply miss Amit. I actually mentored Amit when he first joined Microsoft and then enjoyed all those years as colleagues in engineering. I was really looking forward to partnering with Amit to release Dallas. But I knew the opportunity he was moving to was important for him and his family. I was happy and excited for these folks, but as you can imagine, I quickly was in a few meetings with Asad Khan on how to handle the situation and ensure Dallas moved forward.



For other teams or other companies, losing leaders like this would rock them to the core and greatly affect the ability to release a product. But not SQL Server. Asad was still our PM leader. Peter Carlin, the longtime SQL Server and Azure SQL engineering leader, would become our vice president over all of Azure SQL and SQL Server (and even more today). Rohan Kumar was still our Vice President of Azure Data. Hanuma Kodavalla, Technical Fellow at Microsoft and longtime SQL veteran, would be our engineering leader for all the developers to make Dallas a reality. Asad would promote the famous Joe Sack to run the program managers for all of SQL Server. And in a huge move, we promoted Pedro Lopes to be the technical lead PM for Dallas. Pedro is not only an incredibly talented, well-respected technical leader but someone I personally call a friend. And well, I guess we still had me. SQL Server is in my blood, so of course I needed to stay the course and see project Dallas to the end no matter what challenges came up. In the spring of 2021 with these organizational announcements, we also tentatively agreed on some key aspects to Dallas:

- Project Dallas would be called SQL Server 2022.
- We would announce this release as a Private Preview in the fall of 2021.
- We would announce a Public Preview in late spring of 2022.
- We would target the General Availability of SQL Server 2022 in the fall of 2022.

With a goal for a Private Preview in the fall of 2021 (we made the decision to announce this at the Microsoft Ignite virtual event, which we knew at this time was in early November 2021), we had to get more organized than ever and start cranking out some builds. Joe assigned Pedro to be our de facto release program manager holding the reins to ensure everyone was meeting release criteria and getting features lined up for an announcement.

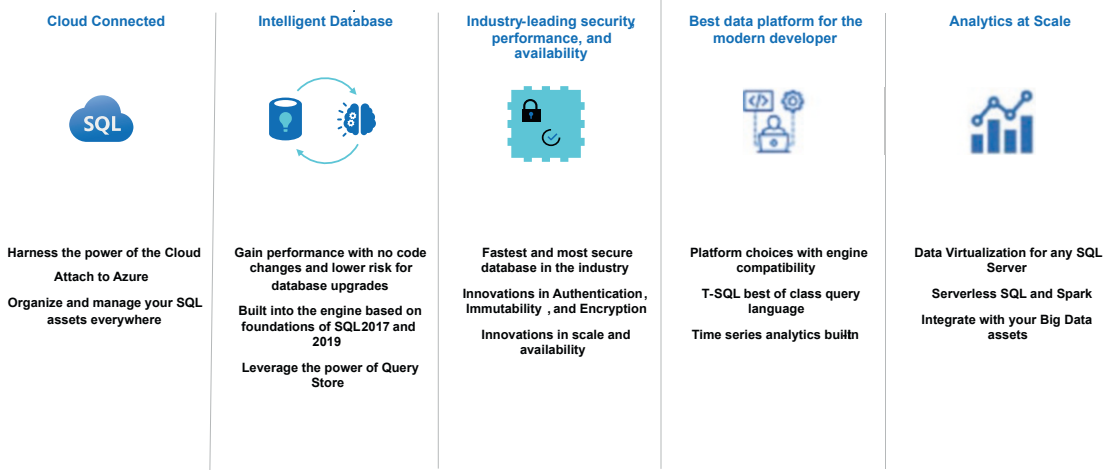
Our marketing team assigned Matthew Burrows to be the lead marketing manager for SQL Server 2022. Matt quickly started building weekly virtual team meetings to talk about SQL vNext (we always call a version vNext until we 100% agree on the exact name) including myself, Joe Sack, Pedro Lopes, James Rowland-Jones, and Sonya Waitman (Sonya works with Matt in marketing and was so valuable throughout the entire release). We would soon add another party to this band, Kendal Van Dyke. Joe hired Kendal to run our Early Adoption Program (EAP) since our first announcement would be a Private Preview program, which required customers to register with Microsoft to test and try out

early builds. Kendal was a longtime veteran of Microsoft as a customer engineer and a well-respected person in the community and within Microsoft. I quickly met with Pedro and Joe and asked how I could best help us march toward Private Preview. They both responded very quickly, “Build us a deck!”

So in the late spring of 2021, I started “doing my thing” to build an “NDA” (for anyone with a nondisclosure agreement with Microsoft) deck we could use. The very first attempt for an overview of SQL Server vNext looked like the following Figure 1-1.

## SQL Server vNext (Project Dallas and Austin)

A hybrid, modern data and analytics platform built on industry-leading security performance, and availability



MICROSOFT CONFIDENTIAL

**Figure 1-1.** The first SQL Server vNext slide

You can see in Figure 1-1 there are not really any specific features listed, just themes. Plus, you will notice at the time we were still debating the future of Big Data Clusters (BDC). Project Austin, named by James Rowland-Jones, was an early proposal to reinvent BDC with serverless SQL and Spark.

---

**Note** You may have read that in February of 2022 we announced the retirement of SQL Server Big Data Clusters, which explains why it was not a part of the SQL Server 2022 release. You can read more at <https://cloudblogs.microsoft.com/sqlserver/2022/02/25/the-path-forward-for-sql-server-analytics>.

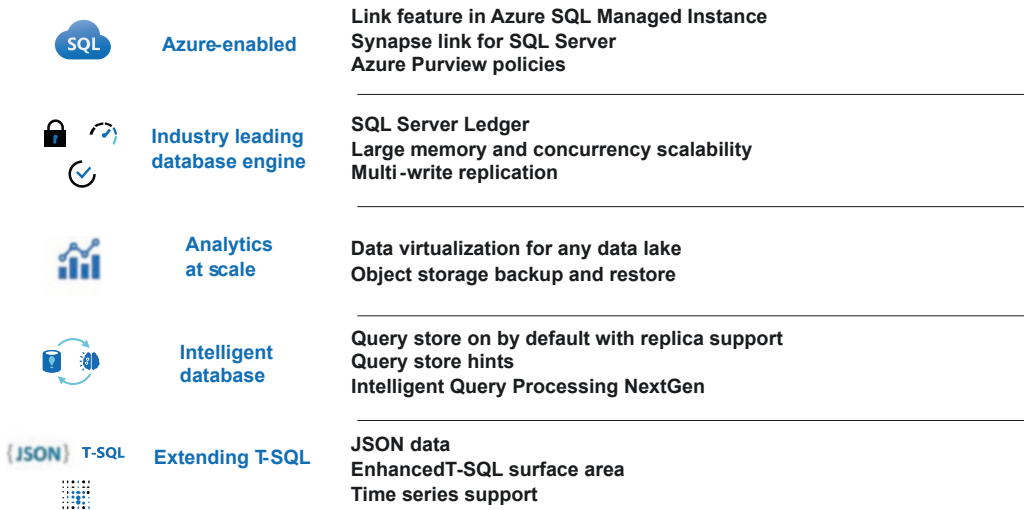
---

I knew all the planned features for Dallas, but at this time I didn't have a 100% commitment that we could talk about all of them yet. One thing I knew is that these "pillars" represented the major categories for features in the release.

To build a foundation for a true public deck on SQL Server vNext, we first needed acknowledgment on the name. We got that from Matthew Burrows about 1 month out from Ignite. SQL Server 2022 would be the name. We also got a much clearer picture of all the features we knew would make the release, so I built this first version of an all-up SQL Server 2022 slide with major features with some rebranding for categories which you can see in the following Figure 1-2.

## SQL Server 2022

A hybrid, data and analytics platform built on industry-leading security performance, and availability



**Figure 1-2.** The first all-up SQL Server 2022 slide

With this in mind, we all decided to keep the announcement for SQL Server 2022 very close to the chest. I was still presenting virtually and even in person (yes, it is true I started presenting at in-person events such as the SQL Server and Azure SQL Conference as early as June 2021). And I would always get asked the question, "When is the next version of SQL Server coming out?" or even better "Will there be a new version of SQL Server?" Both fair questions since we were very quiet on this subject, and by the summer of 2021, it would have been almost 2 years since we released SQL Server or even talked about it.

I then built out more slides talking about each of these areas, including all the features we knew would be in the release. I consulted with Pedro regularly on exactly what we could commit would be in the release, even if it were not immediately in Private Preview builds.

# Announcing SQL Server 2022

Preparing for an announcement of this magnitude is no simple task, and our entire team worked tirelessly in October and right up to the announcement for Ignite to have everything in order.

Asad Khan and Anna Hoffman recorded their Ignite presentation announcing SQL Server 2022 (you can see their original session at <https://learn.microsoft.com/en-us/events/ignite-nov-2021/azure/breakouts/brk225/>).

Joe, Pedro, and Kendal worked to make sure everything was lined up to have the Private Preview build and EAP registration setup.

Matthew and Sonya worked to get our main website (<https://aka.ms/sqlserver2022>) up and running and build a blog, which would be authored by our Corporate Vice President of Azure Data SQL Engineering, Peter Carlin (my manager, Asad Khan, works for Peter, and Peter works for Rohan Kumar) at <https://cloudblogs.microsoft.com/sqlserver/2021/11/02/announcing-sql-server-2022-preview-azure-enabled-with-continued-performance-and-security-innovation>.

As we prepared for the big launch announcement, the famous Microsoft Mechanics team (led by Jeremy Chapman) reached out to me about recording a Mechanics video on SQL Server 2022. Microsoft Mechanics (<https://aka.ms/microsoftmechanics>) is just a cool team. The preparation that goes into the demos, the script, and the animations is a pure marvel. I recorded the session weeks before the announcement in my home in North Richland Hills, Texas, in one take! The result was a very popular video at <https://aka.ms/sqlmechanics22>. It was the #1 watched video produced for Microsoft Ignite. These types of videos are also a forcing function to get demos working even with early builds. I was able to collaborate with the Managed Instance team to get the link feature for Managed Instance working (along with Anna who demonstrated this in Asad's session) with very early builds of both SQL Server 2022 and SSMS (SQL Server Management Studio). I worked with our team to get Synapse Link working and was able to craft up a fun demo (with a big thanks to Anna for building a machine learning model in Spark) with a bit of a Texas twist! And then a demo for what you will probably find out as you read this book is my favorite feature: Parameter-Sensitive Plan (PSP) optimization. PSP is "all engine," so I spun up an Azure Virtual Machine with "pre-CTP" bits and was able to modify a demo that Pedro Lopes and my good friend from development Jack Li gave me.

In the weeks leading up to the announcement, Joe Sack asked me to do a few NDA presentations on SQL Server 2022 to some of our key partners to prepare them for our launch. It was pretty fun to do these as it was the first time I had actually told anyone about SQL Server 2022 outside of Microsoft.

One last piece of promotion for the SQL Server 2022 Private Preview launch was a partnership with Intel for an article in *WIRED* magazine. The resulting article at <https://www.wired.com/sponsored/story/data-is-driving-the-future-of-business> was wildly popular and received a lot of views at launch and for months afterward.

One other interesting twist for me was a trip to Redmond the week before Ignite in late October 2021. Due to the pandemic, I had not traveled to Redmond since February of 2020. I went to Redmond for a few reasons. First, I needed to record a webinar in our Microsoft studios on SQL Server 2022 with Matt Burrows and Anna Hoffman. You can see that webinar on-demand at <https://aka.ms/sqlserver2022webinar>. The second reason was to record a demo for the PASS Summit keynote with Rohan and Peter Carlin. It was fun to record this because I recorded the demo using the PSP demo I did for Microsoft Mechanics, but the twist was having Conor Cunningham interrupt me during my recording. We kind of planned this out, but funny enough we captured it in the first take. It was fun to do since both of us were used to being on stage live with Rohan for the PASS keynote, which couldn't happen in 2021. You can watch the recording of that keynote at <https://youtu.be/Yd1g1KpmrKU?list=PLoGAcXKPCrVYWLmrDZJ9XTdSJAduAefm7>. The third reason I traveled to Redmond was to meet some of my colleagues in person for the first time in almost 2 years. It was great to see folks in person again. While I was up there, Rie Merrit, who runs our MVP program, asked if I would do a preview announcement of SQL Server 2022 for the MVPs since we were making our big announcement the next week. It was really fun to do that session because it was the first time I had told a fairly large audience about SQL Server 2022. I'm used to keeping this type of confidential information, but it is still hard. For example, we had reserved slots for the PASS Summit the week after Ignite for SQL Server 2022 sessions but couldn't even tell the PASS event organizers the details of our sessions. In addition, Thomas LaRock had invited me and others to come speak at their SQL Live event in Orlando in November (<https://sqllive360.com/ECG/live360events/Events/Orlando-2021/SQL.aspx>), and I had reserved a session called "Microsoft SQL Engineering Session Under Wraps." The organizers were asking me, "Can't you tell us anymore about this?" and my answer was "Not really." At the MVP session, I could hear Thomas in his head going, "Ah, that is why Bob couldn't tell us." I had the same issue with sessions I was doing at the SQL Server and Azure SQL Conference (formerly SQLIntersection) in December in Las Vegas.

It was a whirlwind week, which also marked my 28th work anniversary at Microsoft. From SQL OS/2 to now, I would have never dreamed I would be still working on SQL Server.

Finally, on November 2nd at Microsoft Ignite, Scott Guthrie announced SQL Server 2022. We launched all of our other assets along with Asad's session as well at Ignite. I also feel that LinkedIn can be a powerful method to make announcements, so I crafted a LinkedIn article the day of launch at [www.linkedin.com/posts/bobwardms\\_activity-6861322082387664896-mBzK](https://www.linkedin.com/posts/bobwardms_activity-6861322082387664896-mBzK).

I felt the reception of SQL Server 2022 Private Preview was very positive. I think many in the industry and community quite frankly wondered whether we would ever build another major version of SQL Server. I personally saw reviews and comments of surprise that we not only were building another version but also we were including some very fairly major enhancements including all the Azure-connected features, built-in query intelligence, SQL Ledger, and data virtualization. It was a little awkward to make this announcement because we started with Private Preview, which is not how we announced SQL Server 2019. There was a precedence for this with SQL Server 2017, but I realized the frustration this caused for some who were anxious to download the bits and see what was inside.

Right after Ignite was the famous PASS Summit, which was now being organized and run by Redgate, and for 2021 it was all virtual. Microsoft was the primary sponsor, so I lined up four sessions that we at Microsoft would present on SQL Server 2022:

**Introduction to SQL Server 2022** – This was an “all-up” session for everything in SQL Server 2022 that we were announcing for the product.

**SQL Server 2022: The hybrid data platform** – I also did this session as a deeper dive into all of the Azure-connected features of SQL Server 2022.

**SQL Server 2022 storage engine capabilities** – I asked David Pless to do a deep drive into the core engine areas of SQL Server 2022: security, performance, and availability.

**Azure SQL and SQL Server 2022: Intelligent database futures** – I asked Pedro to do a deep drive into the new built-in query intelligence features of SQL Server 2022.