



# Troubleshooting SQL Server Enterprise Geodatabase Performance Issues

Matthew Ziebarth and Ben Lin

# Troubleshooting SQL Server Enterprise Geodatabase Performance Issues

## AGENDA

- General configuration recommendations
- Recovery models & log space
- Trace flag 4199
  
- Max Degree of Parallelism
- Demo on MAXDOP
  
- Final Notes
- Q & A
- 





# General configuration recommendations

# The ArcSDE DBTUNE Table

- Use the DBTUNE table to separate data types
- This is not just for performance, but also for data recovery
- A RAID will not eliminate all disk contention and I/O
- Flash storage will not compensate for data file contention

```
/****** Script for SelectTopNRows command from SSMS *****/
SELECT TOP 1000 [keyword]
      , [parameter_name]
      , [config_string]
FROM [sde103].[dbo].[SDE_dbtune]
```

00 %

	keyword	parameter_name	config_string
1	DATA_DICTIONARY	B_CLUSTER_ROWID	0
2	DATA_DICTIONARY	B_CLUSTER_USER	0
3	DATA_DICTIONARY	B_INDEX_ROWID	WITH FILLFACTOR = 75
4	DATA_DICTIONARY	B_INDEX_USER	WITH FILLFACTOR = 75
5	DATA_DICTIONARY	B_STORAGE	
6	DATA_DICTIONARY	I_STORAGE	
7	DATA_DICTIONARY	MVTABLES_MODIFIED_INDEX	WITH FILLFACTOR = 75
8	DATA_DICTIONARY	MVTABLES_MODIFIED_TABLE	
9	DATA_DICTIONARY	STATE_LINEAGES_INDEX	WITH FILLFACTOR = 75
10	DATA_DICTIONARY	STATE_LINEAGES_TABLE	
11	DATA_DICTIONARY	STATES_INDEX	WITH FILLFACTOR = 75
12	DATA_DICTIONARY	STATES_TABLE	
13	DATA_DICTIONARY	VERSIONS_INDEX	WITH FILLFACTOR = 75
14	DATA_DICTIONARY	VERSIONS_TABLE	
15	DATA_DICTIONARY	XML_INDEX_TAGS_INDEX	WITH FILLFACTOR = 75
16	DATA_DICTIONARY	XML_INDEX_TAGS_TABLE	
17	DEFAULTS	A_CLUSTER_RASTER	0
18	DEFAULTS	A_CLUSTER_ROWID	0
19	DEFAULTS	A_CLUSTER_SHAPE	1
20	DEFAULTS	A_CLUSTER_STATEID	0
21	DEFAULTS	A_CLUSTER_USER	0
22	DEFAULTS	A_CLUSTER_XML	0
23	DEFAULTS	A_INDEX_RASTER	WITH FILLFACTOR = 75
24	DEFAULTS	A_INDEX_ROWID	WITH FILLFACTOR = 75
25	DEFAULTS	A_INDEX_SHAPE	WITH FILLFACTOR = 75
26	DEFAULTS	A_INDEX_STATEID	WITH FILLFACTOR = 75
27	DEFAULTS	A_INDEX_USER	WITH FILLFACTOR = 75
28	DEFAULTS	A_INDEX_XML	WITH FILLFACTOR = 75
29	DEFAULTS	A_MS_SPINDEX	GRIDS = (MEDIUM, M...
30	DEFAULTS	A_OUT_OF_ROW	0
31	DEFAULTS	A_PARTITION_RASTER	
32	DEFAULTS	A_PARTITION_ROWID	
33	DEFAULTS	A_PARTITION_SHAPE	
34	DEFAULTS	A_PARTITION_STATEID	
35	DEFAULTS	A_PARTITION_USER	
36	DEFAULTS	A_PARTITION_XML	
37	DEFAULTS	A_STORAGE	

# Types of data to separate

- In DBTUNE:
  - Business tables from delta tables
  - Business indexes from delta indexes
  - Complex dataset types such as networks, topologies
- On DISK:
  - Data files from log files





# Recovery models & log space

# The log may be slowing you down

- LDF files become very large
- Writing to any log file takes time



# Recovery Models

- **Simple Model.** No log backups. Work-loss exposure = last full database backup. This may be suitable for your organization.
- **Full Model.** Large log backups including all transactions. Near-zero work-loss exposure. These logs must be stored on server as well as backed up to ensure full historical point in time recovery. Data file corruption, no problem.
- **Bulk Model.** Reduced size of Full but still large log files. Work-loss exposure = last log backup. No point in time recovery.



+





# Recovery Models

The screenshot displays the 'Database Properties - master' window in SQL Server Enterprise Manager. The 'Options' page is selected in the left-hand navigation pane. The main area shows several configuration options for the database, with the 'Recovery model' set to 'Simple'. Below these options is a table of 'Other options'.

Script Help

Collation: SQL\_Latin1\_General\_CP1\_CI\_AS

Recovery model: Simple

Compatibility level: SQL Server 2012 (110)

Containment type: None

Other options:

Auto Shrink	False	^
Auto Update Statistics	True	
Auto Update Statistics Asynchronously	False	

# Log Location

Database Properties - master

Select a page

- General
- Files
- Filegroups
- Options
- Change Tracking
- Permissions
- Extended Properties

Script Help

Database name: master

Owner: sa

Use full-text indexing

Database files:

Logical Name	File Type	Filegroup	Initial Size (MB)	Autogrowth / Maxsize
master	Rows ...	PRIMARY	4	By 10 percent, Unlimited
mastlog	Log	Not Applicable	1	By 10 percent, Unlimited



# Trace flag 4199

“The Principle of Least Surprises”

# Trace Flag 4199

- Trace flag 4199, per Microsoft:

**“Controls multiple query optimizer changes previously made under multiple trace flags. “**

**And...**

**“...any hotfix that could potentially affect the execution plan of a query must be controlled by a trace flag. Except for fixes to bugs that can cause incorrect results or corruption, these hotfixes are turned off by default, and a trace flag is required to enable the fix.”**

- **Enable this globally to take advantage of bug fixes.**



# Max Degree of Parallelism



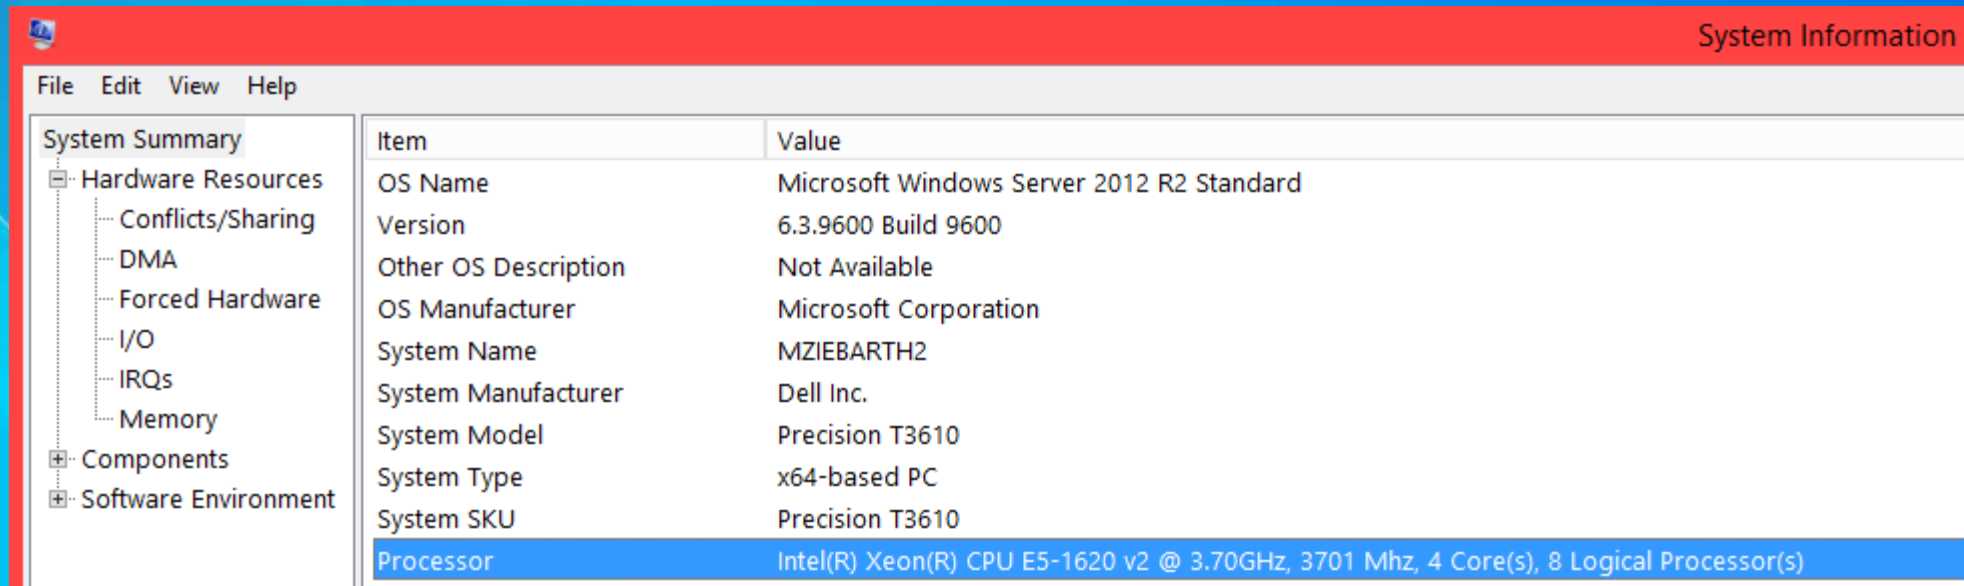
# What comprises Degree of Parallelism?

- The DOP is based on multiple factors of the OS and the Database Instance:
  - # of available CPU Cores
  - # of available logical threads
  - Type of query
  - # of rows returned
  - Availability of distribution statistics
- **Hardware + Utilization + Query + Output + Stats = DOP**



# What is a CPU Core Count?

- The physical number of cores on the CPU within the processor packages
- To find this: Click on Run and type MSINFO32

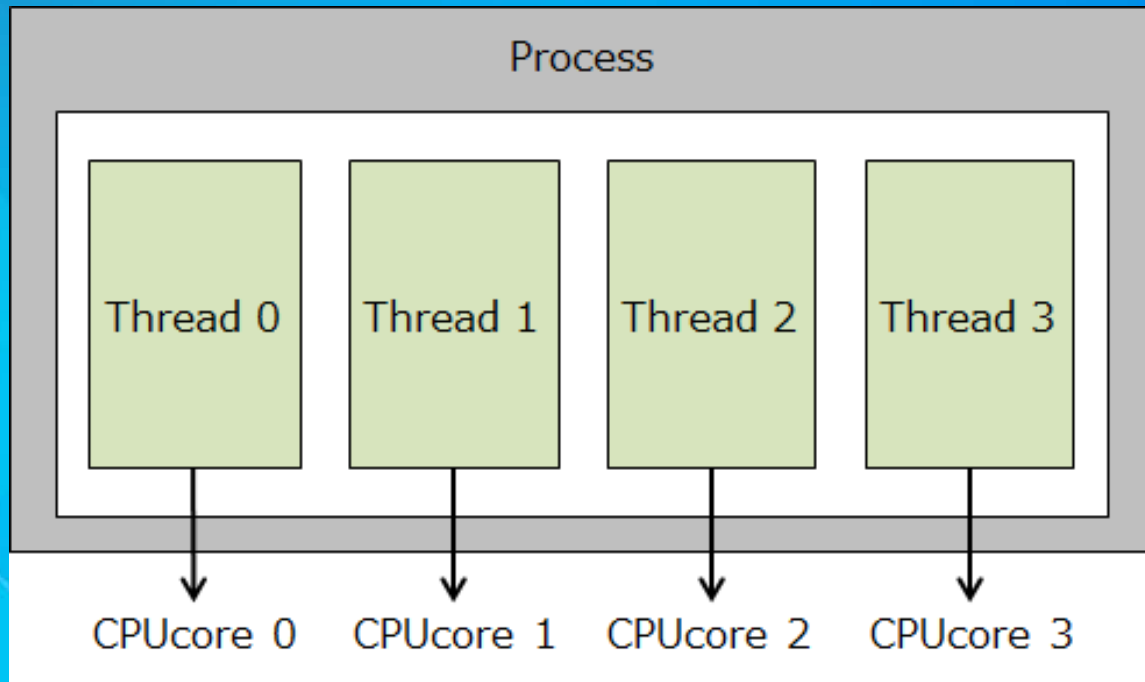


The screenshot shows the Windows System Information utility. The window title is "System Information". The left sidebar shows a tree view with "System Summary" selected. The main pane displays a table of system information. The "Processor" row is highlighted in blue.

Item	Value
OS Name	Microsoft Windows Server 2012 R2 Standard
Version	6.3.9600 Build 9600
Other OS Description	Not Available
OS Manufacturer	Microsoft Corporation
System Name	MZIEBARTH2
System Manufacturer	Dell Inc.
System Model	Precision T3610
System Type	x64-based PC
System SKU	Precision T3610
Processor	Intel(R) Xeon(R) CPU E5-1620 v2 @ 3.70GHz, 3701 Mhz, 4 Core(s), 8 Logical Processor(s)

# What is the Thread Count?

If the CPU uses multi-threading, the logical thread count is the core count x2.



x 2



# What Type of Query?

- The query operators, or query steps, need to be expensive to take advantage of parallelism. Otherwise the query will execute serially.
- Mathematically, parallelism is used when the estimated cost of the operation is higher than the **cost threshold for parallelism** setting.

Results		Messages			
	name	minimum	maximum	config_value	run_value
1	cost threshold for parallelism	0	32767	5	5

# What are Distribution Stats?

- Refers to high or low (poor) selectivity
- Unique key has high selectivity; there is 1

- Ben will demonstrate this in the DEMO

DBCC SHOW_STATISTICS (Itens, Stats_Quantidade) WITH HISTOGRAM					
	RANGE_HI_KEY	RANGE_ROWS	EQ_ROWS	DISTINCT_RANGE_ROWS	AVG_RANGE_ROWS
1	NULL	0	513	0	1
2	100	0	56	0	1
3	104	171	59	3	57
4	107	88	60	2	44
5	111	160	64	3	53,33333
6	118	304	60	6	50,66667

RANGE\_HI\_KEY = Key value from each histogram step  
Ex: In the step of the row 5 we have the value 111 that goes from 108 (107 (row 4) + 1) to 111

RANGE\_ROWS = Number of rows that have values equal to the range excluding the value of RANGE\_HI\_KEY  
Ex: The step of row 5 goes from 108 to 110 (excluding the value 111(RANGE\_HI\_KEY)). There are 160 rows on this range


EQ\_ROWS = Number of rows in the key value of the step  
Ex: There are 64 rows for the value 111 (row 5)

DISTINCT\_RANGE\_ROWS = Number of distinct values within the range. Excluding the value of RANGE\_HI\_KEY.  
Ex: There are 3 distinct values in the step of the row 5 (108 to 110)  
Should call DISTINCT\_RANGE\_VALUES instead of DISTINCT\_RANGE\_ROWS

AVG\_RANGE\_ROWS = Average of rows for values within the range (RANGE\_ROWS / DISTINCT\_RANGE\_ROWS)  
Ex: The average of rows for each value within the step of the row 5 is 53,33333 (160/3)

The actual selectivity is not as important for DOP as maintaining these statistics.

# The Cost Threshold for Parallelism

- If the SQL Server query optimizer believes a query operator will take this amount of cost units (in seconds) if run serially, this parameter is reached.
- The default is 5 and can be set between 0 and 32767 seconds, or about 9 hours.
- **Threshold is reached, therefore parallelism!** 
- This is a very old default value but your data design factors into changing it.

**The coordination of the parallel threads also has a hidden cost!**

**Does this affect you?**



# Why is this happening?

- It's a complex bug in that the inverse of what you would expect to occur, occurs.
- Query Optimizer incorrectly computes it would be fastest to coordinate all available threads to return your spatial results.

The Spatial Index is not used...



- This is intermittent because if utilization is high and spare cores do not exist to use within a query, parallelism isn't an option sometimes.

# How do we diagnose?

- This will come in to your GIS/T department as a performance problem that will be intermittent and inverse to server load.
- Get the full configuration of the SQL Server database (exact build) and the database server hardware such as CPUs, Cores, Processors, Model/Brand, and RAM.
- Ask about utilization trends.
- SQL trace or Execution plan check for spatial index.



**\*This issue only applies to SQL Server Geometry or Geography types.**

# How do we fix this?

- The fastest way to diagnose and “fix” this is to alter the MAXDOP setting.  
(In a Test environment first, preferably)
- 1 = DISABLE parallelism. 2 = 2 cores per operator, and so on.
- Generally users should find a “sweet spot” MAXDOP setting or adjust per data.

**THIS SETTING DOES NOT REQUIRE A RESTART OF ANY KIND. IT IS DYNAMIC.**

- When you can, patch this. Please see KB articles on final page.

**\*The default was finally changed from 0 to a hardware calculation in SS 2016! 😊**



# MAXDOP Demo

Ben Lin



# Final Notes

- Check your MAXDOP settings!
- Make sure your feature classes have a spatial index
- Use Compress, Analyze Datasets, Rebuild Indexes often
- Use SQL Server Profiler
- Examine execution plans
- Enable SDE Intercept
- Contact Esri Technical Support for assistance

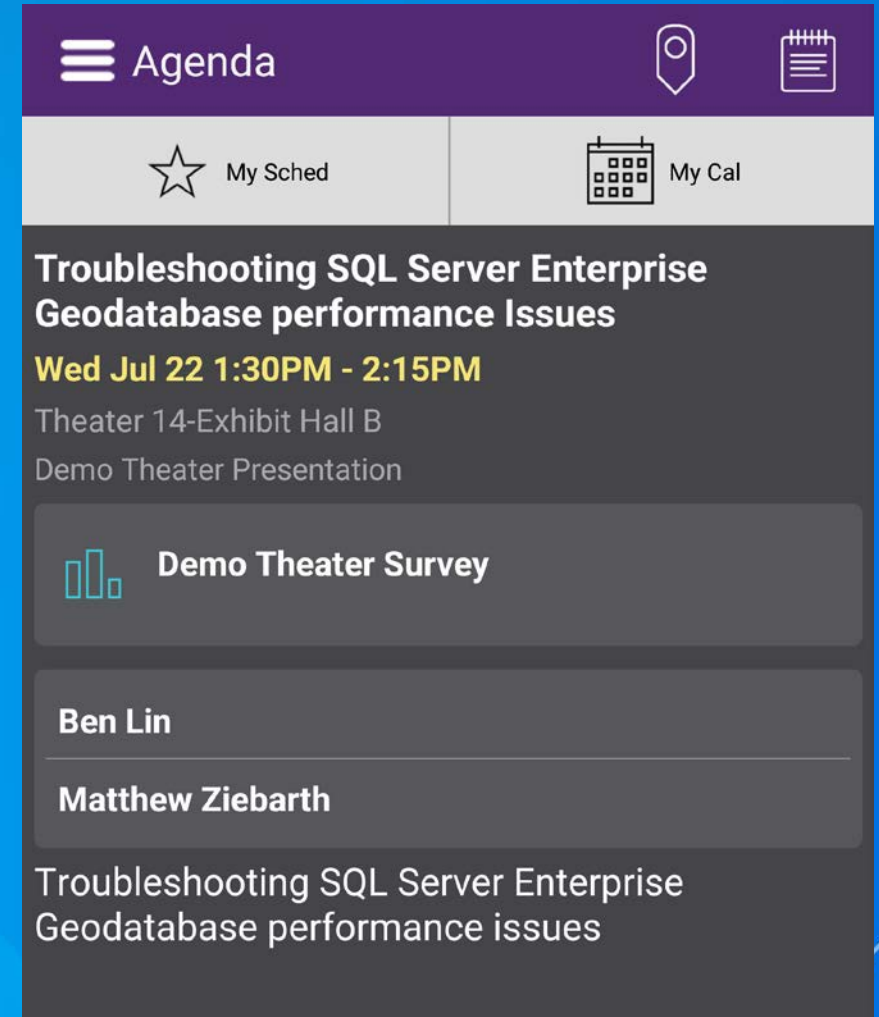


```
RPC:Completed declare @p1 int set @p1=NULL exec... SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac
RPC:Completed declare @p1 int set @p1=1073741842... SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac
RPC:Completed exec sp_cursorfetch 180150035,2,1,100 SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac
RPC:Completed exec sp_cursorclose 180150035 SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac
RPC:Completed exec sp_cursorunprepare 1073741842 SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac
RPC:Completed declare @p1 int set @p1=5 exec sp... SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac
RPC:Completed exec sp_unprepare 5 SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac
RPC:Completed declare @p1 int set @p1=6 exec sp... SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac
RPC:Completed exec sp_unprepare 6 SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac
RPC:Completed declare @p1 int set @p1=7 exec sp... SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac
RPC:Completed exec sp_unprepare 7 SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac
RPC:Completed declare @p1 int set @p1=8 exec sp... SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac
RPC:Completed exec sp_unprepare 8 SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac
RPC:Completed declare @p1 int set @p1=9 exec sp... SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac
RPC:Completed exec sp_unprepare 9 SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac
SQL:BatchStarting set fmtonly on EXEC geoprod.dbo.S... SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac
SQL:BatchCompleted set fmtonly on EXEC geoprod.dbo.S... SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac
SQL:BatchStarting set fmtonly off SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac
SQL:BatchCompleted set fmtonly off SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac
RPC:Completed declare @p1 int set @p1=10 exec s... SDE:8248:9060 geoprod_SQLsysadmin CITYOFPCLOVIS\geoprod_SQLsysac

declare @p1 int
set @p1=1073741841
declare @p2 int
set @p2=180150033
declare @p5 int
set @p5=1
declare @p6 int
set @p6=1
declare @p7 int
set @p7=0
exec sp_cursorprepexec @p1 output @p2 output,NULL,N'SELECT logfile_name, logfile_id, logfile_data_id, registration_id, flags, session_tag,
column_name, logfile_data_db, logfile_data_owner, logfile_data_table FROM geoprod.DBO.SDE_logfiles WHERE flags%2 = 0',@p5 output @p6 output @p7 output
select @p1, @p2, @p5, @p6, @p7
```

# Thank you...

- Please fill out the session survey in your mobile app
- Select “Troubleshooting SQL Server Enterprise Geodatabase Performance Issues” in the Mobile App
- Use the Search Feature to quickly find this title
- Click “Demo Theater Survey”
- Answer a few short questions and enter any comments



# Want to learn more?

- **Esri Documentation**
- [ArcGIS 10.3 SQL Server DBTUNE configuration parameters](#)
- [Esri KB 35704 - SDE Intercept](#)
- [Esri KB 38871 - SQL Server MAXDOP](#)
- [Esri KB 36617 - SQL Server Spatial Index Hinting](#)
- **Microsoft Documentation**
- [MS KB 2570501 - MAXDOP Bug Fix for 2008](#)
- [MS KB 2757097 - MAXDOP bug in 2012 - Defect 630366](#)
- [Defect 20179508 - Spatial index not used on multiproc machine unless MAXDOP set downward](#)
- [MS Premier Field Engineer MAXDOP tips](#)
- [SQL Server Recovery Models](#)
- [Trace flag 4199](#)
- **Give us a call at ESRI Technical Support at 888-377-4575**

The background features a vibrant blue gradient. On the left side, there are several overlapping geometric shapes: a large purple triangle pointing upwards, a yellow triangle pointing downwards, and a dark purple triangle pointing downwards. The yellow triangle contains a faint, light-colored map of a region, possibly a county or state map. The overall composition is modern and abstract.

Questions?



Understanding our world.