

The Data Repository Overnight DBA

2018 MUSE International

Educational Session #1104

Thursday, May 31 2:45pm

Presenter: Ian Proffer

Today's agenda



What is a DBA?

SQL Server & Management Studio

MEDITECH's expectations for you

Administration: database maintenance, managing growth, user access and security

Analysis: getting data efficiently, monitoring and optimizing performance



What is a “DBA?”

- A. A database administrator
- B. A database analyst
- C. A database architect
- D. An FTE that hasn't been filled yet
- E. A and B (and maybe D)



Working with tables in DR

The screenshot displays the Microsoft SQL Server Management Studio (SSMS) interface. The 'Object Explorer' pane on the left shows the server hierarchy for 'SURFACEIP (SQL Server 14.0.1000.169 - SURFACEIP\ianpr)'. The 'Tables' folder is selected, and the 'Object Explorer Details' pane on the right displays a list of tables in the 'traindb' database. The table list includes columns for Name, Schema, Create Date, Data Space Used (KB), and Row Count. The 'Tables' folder summary at the bottom indicates 96 items.

Name	Schema	Create Date	Data Space Used (KB)	Row Count
AbsDrgData	dbo	9/14/2017 4:02 PM	160	599
AbsDrgDiagnoses	dbo	9/14/2017 4:02 PM	48	506
AbsDrgProcedures	dbo	9/14/2017 4:02 PM	8	33
AbsInsurances	dbo	9/14/2017 4:02 PM	96	533
AbstractData	dbo	9/14/2017 4:02 PM	344	920
AbstractProviders	dbo	9/14/2017 4:02 PM	48	722
AdmDischarge	dbo	9/14/2017 4:02 PM	72	418
AdmEmployers	dbo	9/14/2017 4:02 PM	32	462
AdmInsuranceOrder	dbo	9/14/2017 4:02 PM	32	700
AdmInsuredInfo	dbo	9/14/2017 4:02 PM	64	544
AdmittingData	dbo	9/14/2017 4:02 PM	128	1440
AdmNursingCensus	dbo	9/14/2017 4:02 PM	8768	100115
AdmProviders	dbo	9/14/2017 4:02 PM	72	1465
AdmVisitClinicalQueries	dbo	9/14/2017 4:02 PM	3480	34282
AdmVisitClinicalQueriesMult	dbo	9/14/2017 4:02 PM	848	13635
AdmVisitEvents	dbo	9/14/2017 4:02 PM	736	3395
AdmVisits	dbo	9/14/2017 4:02 PM	464	1447
BarBchs	dbo	9/14/2017 4:02 PM	1144	2847
BarBchTxns	dbo	9/14/2017 4:02 PM	11528	69575
BarBills	dbo	9/14/2017 4:02 PM	200	696
BarChargeTransactions	dbo	9/14/2017 4:02 PM	16800	93489
BarCollectionTransactions	dbo	9/14/2017 4:02 PM	16	51
BarEmployers	dbo	9/14/2017 4:02 PM	40	462
BarInsuranceLedger	dbo	9/14/2017 4:02 PM	400	2182
BarInsuranceOrder	dbo	9/14/2017 4:02 PM	24	523
BarInsurances	dbo	9/14/2017 4:02 PM	184	655
BarInsuredData	dbo	9/14/2017 4:02 PM	72	589
BarVisitFinancialData	dbo	9/14/2017 4:02 PM	504	1571
BarVisitGuarantors	dbo	9/14/2017 4:02 PM	88	470
BarVisitProviders	dbo	9/14/2017 4:02 PM	200	2100

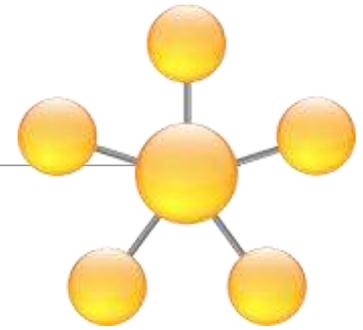
Using Management Studio



Manage your Microsoft SQL Server instances, including:

- System-wide configuration and security, scheduled job execution
- Create and maintain databases, including security administration (i.e. DBA activities)
- Analyze tables and data
- Write queries to get report results
- Create database programming objects – stored procedures

SQL Server & Data Repository



- DR is not a typical database (it's a relational database without relationships)
- C/S and MAGIC 5.67 have 8,000+ tables and 80,000+ data fields
- Expanse/6.1 has 19,000+ tables and 120,000+ data fields in 2 databases
- Security is managed through SQL Server or Windows Active Directory (no connection to MEDITECH security mechanisms)
- DR databases include tables with primary key clustered indexes and stored procedures, but no defined foreign key relationships

MEDITECH Wants You!



To take care of your
Data Repository



Routine Maintenance Recommendations

SQL Server

Monitor SQL Server error and Windows Event logs

- It is the responsibility of the customer to resolve SQL Server and Windows-specific issues. The customer should work with the vendor to resolve these errors. As stated previously, if necessary, MEDITECH will work with you to direct you to the appropriate resources or whether they do not impact your business.
- Errors noted in the SQL Server error log and Windows Event logs that may affect DR should be reported to MEDITECH.

Performance Monitors

- A consistent pattern of ever-increasing CPU and Memory utilization may be indicative of an inefficient use of resources. MEDITECH is closely integrated with Microsoft's Performance Monitors to determine if issues are the result of inefficient use of Microsoft SQL software, or other vendor software. MEDITECH may then direct you to the appropriate hardware needs upgrading. Any initial performance issues should be reported to MEDITECH and if necessary to your hardware supplier.

Database Backup

- MEDITECH recommends the livedb database and/or transaction log be backed up on a daily basis (Please see [SQL Database Backup](#))

Weekly Monitoring

SQL Server

Monitor Space Usage

- Since Data Repository does not automatically purge data and the SQL databases will continue to grow (based on data volume), database expansion recommendations can be found in [Knowledge Base article 30790](#). If SQL databases should be done regularly. Database expansion recommendations can be found in [Knowledge Base article 30790](#). If SQL databases are not able to expand appropriately, DR Transfers will be stopped until additional space is acquired/all appropriate tables within the database.
- Storage History Routines can also be utilized to assist with monitoring database growth and determining when more space is needed. See [30790](#) for additional recommendations for analyzing database growth.
- Additionally, the stored procedure 'sp_spaceused' may be useful in monitoring the space usage of the specified object.

Monitor Index Fragmentation

- As records are added & removed from tables, the indexes that organize that data gradually "fragment", negatively impacting performance. Both clustered and non-clustered indexes should be monitored regularly to identify any fragmentation. MEDITECH has [additional information](#) on the topic of fragmentation.

When two disparate systems (e.g., Data Repository and SQL Server) are used, it is important that any problems are reported to the appropriate vendor.

It is important that any problems are reported to the appropriate vendor.

Microsoft SQL Server is closely integrated with Microsoft's Performance Monitors to determine if issues are the result of inefficient use of Microsoft SQL software, or other vendor software. MEDITECH may then direct you to the appropriate hardware needs upgrading.

Microsoft SQL Server is closely integrated with Microsoft's Performance Monitors to determine if issues are the result of inefficient use of Microsoft SQL software, or other vendor software. MEDITECH may then direct you to the appropriate hardware needs upgrading.

If any reason will appear in red. If you're unable to resolve the issue, please contact your account manager.

When the logs currently being processed. If you are unable to resolve the issue, please contact your account manager.

When no pending activity to transfer they are in progress. If you're unable to resolve the issue, please contact your account manager. Transfer status hasn't been updated for all transfers. Transfer/Update Transfer Routines can be used to update the Data Repository transfer jobs.

Server and database administration

Server monitoring & maintenance

- Check the Windows Server and SQL Server logs regularly or when you see symptoms of possible problems
- Perform regular database maintenance, including backups and data integrity checking
- Monitor overall disk space and usage



Backups and data integrity

Database backups

- Ideally every database has a full backup done daily during non-business hours
 - Older hardware with very large databases may use differential backups
- Do I need to back up the transaction log?
 - The “Simple” recovery model
 - In the event of a system failure, how quickly can the database be brought back to a given point in time?



Data integrity checking (DBCC)

- Check data integrity regularly (weekly, monthly)
 - DBCC, SQL Maintenance Wizard
- Analyze and address index fragmentation

Managing disk and database space



glass = database
water = your data
coaster = disk drive



Checking database & file space

1. Management Studio Report

Disk Usage

[Livefdb]

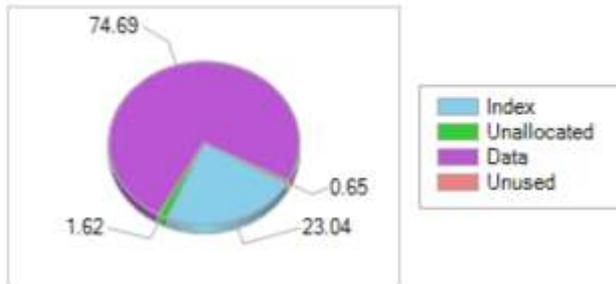
on NMN-DR01 at 5/9/2018 12:32:43 PM



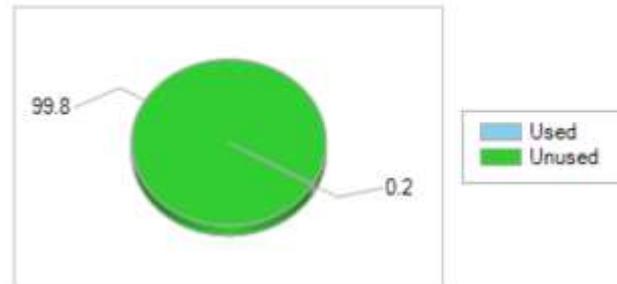
This report provides overview of the utilization of disk space within the Database.

Total Space Usage:	1,072,693.69 MB
Data Files Space Usage:	1,056,627.25 MB
Transaction Log Space Usage:	16,066.44 MB

Data Files Space Usage (%)



Transaction Log Space Usage (%)



Data/Log Files Autogrow/Autoshrink Events

Event	Logical File Name	Start Time	Duration (ms.)	Change In Size (MB)
Data File Auto Growth	Livefdb	5/7/2018 2:33:41 PM	48526	20,718.19

Disk Space Used by Data Files

Filegroup Name	Logical File Name	Physical File Name	Space Reserved	Space Used
PRIMARY	Livefdb	E:\MEDIFILES\Livefdb.mdf	1031.86 GB	1015.40 GB



Checking database & file space

2. Database Properties

The image shows two overlapping windows from SQL Server Enterprise Manager. The top window displays the 'Backup' section of the 'Database Properties - Livedb' dialog, showing the last backup times. The bottom window displays the 'General' section of the same dialog, showing database metadata and file information.

Database Properties - Livedb (Backup Section)

Backup Type	Last Backup Time
Last Database Backup	5/9/2018 12:09:59 PM
Last Database Log Backup	None

Database Properties - Livedb (General Section)

Database: Livedb
Owner: NT AUTHORITY\SYSTEM
 Use full-text indexing

Database files:

Logical Name	File Type	Filegroup	Initial Size (MB)	Autogrowth / Max...	Path
Livedb	Rows Data	PRIMARY	1056628	By 2 percent, ...	E:\Medi...
Livedb_log	Log	Not Applicable	16067	By 2 percent, ...	D:\MED

Checking database & file space

3. T-SQL sp_spaceused

```
SQLQuery1.sql - NM...BS\iproffer (157))* X  
USE Livefdb  
GO  
EXEC sp_spaceused  
GO
```

100 %

Results Messages

	database_name	database_size	unallo
1	Livefdb	1072693.69 MB	17067

	reserved	data	index_
1	1064508752 KB	771125792 KB	28631

```
SQLQuery1.sql - NM...BS\iproffer (157))* X  
USE Livefdb  
GO  
EXEC sp_spaceused 'RegAcct_Main'  
GO
```

100 %

Results Messages

	name	rows	reserved	data	index_size	unused
1	RegAcct_Main	2880018	2002352 KB	857896 KB	1142600 KB	1856 KB



Checking database & file space

4. Query the system tables directly

```
SQLQuery1.sql - NM...BS\iproffer (157))* ×  
  
SELECT  
    file_id, name, type_desc,  
    physical_name, size, max_size  
FROM sys.database_files
```

100 %

Results Messages

	file_id	name	type_de...	physical_name	size	max_size
1	1	Livfdb	ROWS	E:\Medifiles\Livfdb.mdf	135248288	-1
2	2	Livfdb_log	LOG	D:\MEDIFILES\Livfdb_log.ldf	2056504	268435456

Managing disk and database space

Disk space vs. database space

- Make sure primary data drive has plenty of capacity for live database growth
- MT recommends you keep 25% free space for the live databases (typically the E: drive)
- Database auto-growth options
 - Percentage vs. amount of space
 - Unrestricted vs. restricted space



What happens when the drive is full?

- Transfers stop (crash) until the db can grow
- Avoid this at all costs!

Monitoring growth over time

```
SELECT
    YEAR(DateTimeID) AS Year,
    MONTH(DateTimeID) AS Month,
    SUM(TableDataPages) AS DataPages,
    SUM(TableIndexPages) AS IndexPages
FROM
    Livefdb.dbo.DR_SysStorageHistory ←
WHERE
    TableNameID = 'RegAcct_Main'
GROUP BY
    YEAR(DateTimeID),
    MONTH(DateTimeID)
ORDER BY
    Year DESC, Month DESC
```

70 %

Results Messages

	Year	Month	DataPages	IndexPages
2	2018	4	25510032	33160720
3	2018	3	26295000	33428304
4	2018	2	23544280	30473976
5	2018	1	26676568	35840600
6	2017	12	26636200	36137528
7	2017	11	25558920	34506280

DR_SysStorageHistory can be queried to summarize data space usage by date and table. (*SysStorage* in MAGIC and C/S.)

This is helpful if you want to predict future database growth based on the last year of transfer activity.



Managing security & access



Create your own database for report objects (stored procedures, functions, views, etc.)

- This avoids co-mingling your code and objects with Meditech's
- This allows you to limit direct access to live databases

Use Active Directory groups for easier administration

- e.g. **ReportAdmins** group, **ReportViewer** group

For those that need access to live data consider:

- Using views in your database to selectively expose data
- Using SQL database-level roles for direct table access
 - **ReportAdmins**: db_owner, db_denydatawriter (read all data from livedb, ability to create indexes but no INSERT, UPDATE or DELETE)
 - **ReportViewers**: db_datareader (SELECT from tables)
 - **ReportViewers_NoPP**: db_datareader with further DENY SELECT permissions*



An example of a DENY script

```
SELECT 'DENY SELECT ON ' + name + ' TO ReportViewers_NoPP'  
FROM livedb.dbo.sysobjects  
WHERE xtype = 'U' AND name LIKE 'Pp%'  
ORDER BY name
```

	(No column name)
1	DENY SELECT ON PpApplicantComments TO ReportViewers_NoPP
2	DENY SELECT ON PpApplicantEducationDetails TO ReportViewers_NoPP
3	DENY SELECT ON PpApplicantEmployeeAddDetails TO ReportViewers_NoPP
4	DENY SELECT ON PpApplicantEmployeeAddresses TO ReportViewers_NoPP
5	DENY SELECT ON PpApplicantEmployeeDetails TO ReportViewers_NoPP
6	DENY SELECT ON PpApplicantEmployeeDocuments TO ReportViewers_NoPP

Database analysis

Optimizing queries

- Query execution plans

Creating your own indexes

Handling index fragmentation

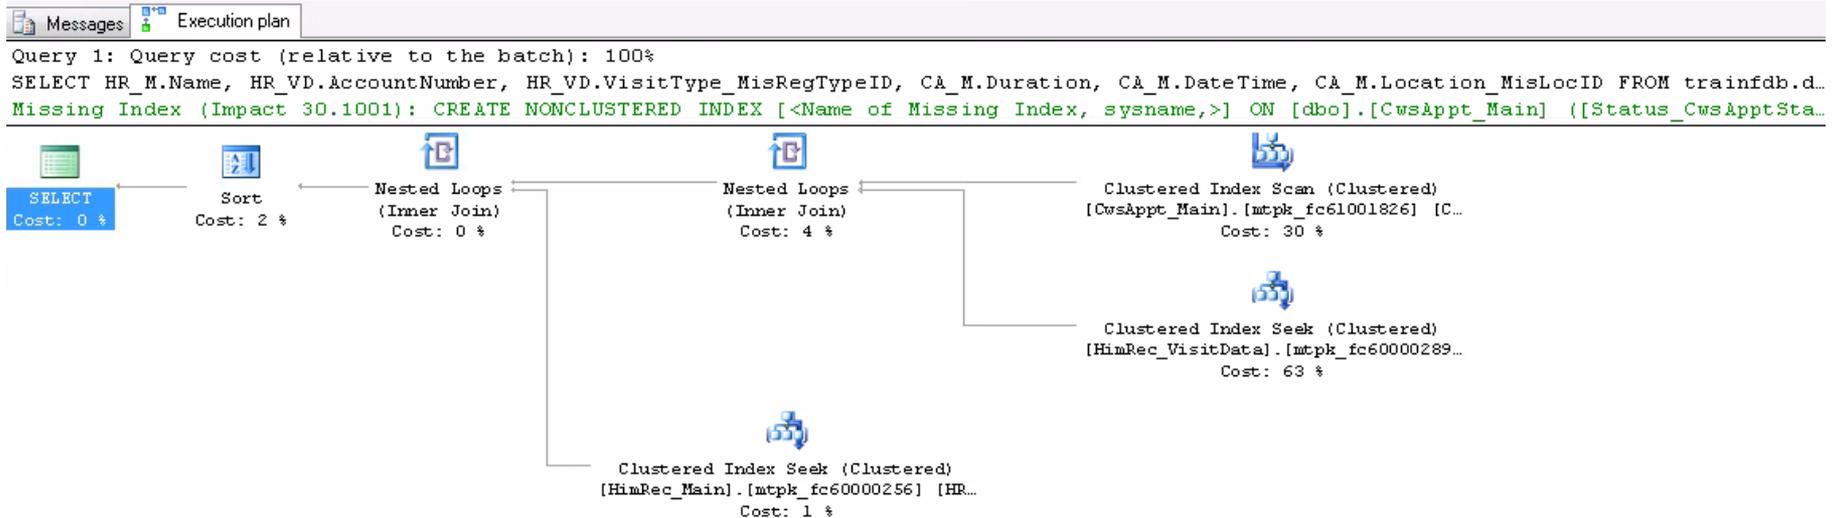
Other performance tools



Optimizing queries

Viewing the query execution plan

```
SELECT
  HR_M.Name,
  HR_VD.AccountNumber,
  HR_VD.VisitType_MisRegTypeID,
  CA_M.Duration,
  CA_M.DateTime,
  CA_M.Location_MisLocID
FROM trainfdb.dbo.HimRec_Main HR_M
INNER JOIN trainfdb.dbo.HimRec_VisitData HR_VD
ON HR_M.SourceID = HR_VD.SourceID
```



Creating non-clustered indexes

Clustered vs. non-clustered

- All DR tables already have a primary key clustered index. Don't change them!
- Non-clustered indexes can be added wherever you like.



Document and save scripts for your indexes

General guidelines for new non-clustered indexes:

- SourceID, VisitID and SourceID, PatientID in any top-level table where it's not already a primary key
 - AbstractData, BarVisits, LabSpecimens, PhaRx, OeOrders, OmOrd_Main
- Fields with highly discrete data values that are used as report parameters (e.g. DateTime)

Fewer general purpose indexes are better than many purpose-built indexes (indexes consume space and require maintenance)

Index fragmentation

The SQL Maintenance Wizard can address this, but...

- Every table in the database is analyzed and acted upon
- Indexes may be dropped and recreated

Doing your own, targeted index maintenance is better

- Assemble a list of tables you use regularly for reports
- If maintenance hasn't been done, address tables selectively on a series of first passes, then on a regular schedule (weekly, monthly, after initial loads)
- After important tables are done, address the others

Limit amount of work done based on a fragmentation threshold or set number of tables (based on size) per session



Analyzing & correcting index fragmentation

First evaluate it with a system dynamic management view

```
use Livefdb
go
select
    index_id, index_type_desc, avg_fragmentation_in_percent,
    fragment_count, avg_fragment_size_in_pages, page_count
from sys.dm_db_index_physical_stats |
(DB_ID(N'Livefdb'), OBJECT_ID(N'RegAcct_Main'), NULL, NULL , NULL);
```

70 %

Results Messages

	index_id	index_type_desc	avg_fragmentation_in_percent	fragment_count	avg_fragment_size_in_pages	page_count
1	1	CLUSTERED INDEX	99.9011673441987	107163	1.00083051053069	107252
2	20	NONCLUSTERED INDEX	99.906611878969	16050	1.0007476635514	16062
3	21	NONCLUSTERED INDEX	99.6913770754892	16155	1.00284741566079	16201
4	22	NONCLUSTERED INDEX	99.4708994708995	15991	1.0046276030267	16065
5	23	NONCLUSTERED INDEX	99.9003611906838	16048	1.00062313060818	16058
6	24	NONCLUSTERED INDEX	99.9439182452642	16043	1.00031166240728	16048
7	26	NONCLUSTERED INDEX	99.9866372686577	14967	1	14967
8	29	NONCLUSTERED INDEX	99.8637992831541	13934	1.00114827041768	13950
9	30	NONCLUSTERED INDEX	99.5087941689114	31555	1	31555



Analyzing & correcting index fragmentation

Then correct it. Indexes can be either reorganized or rebuilt:

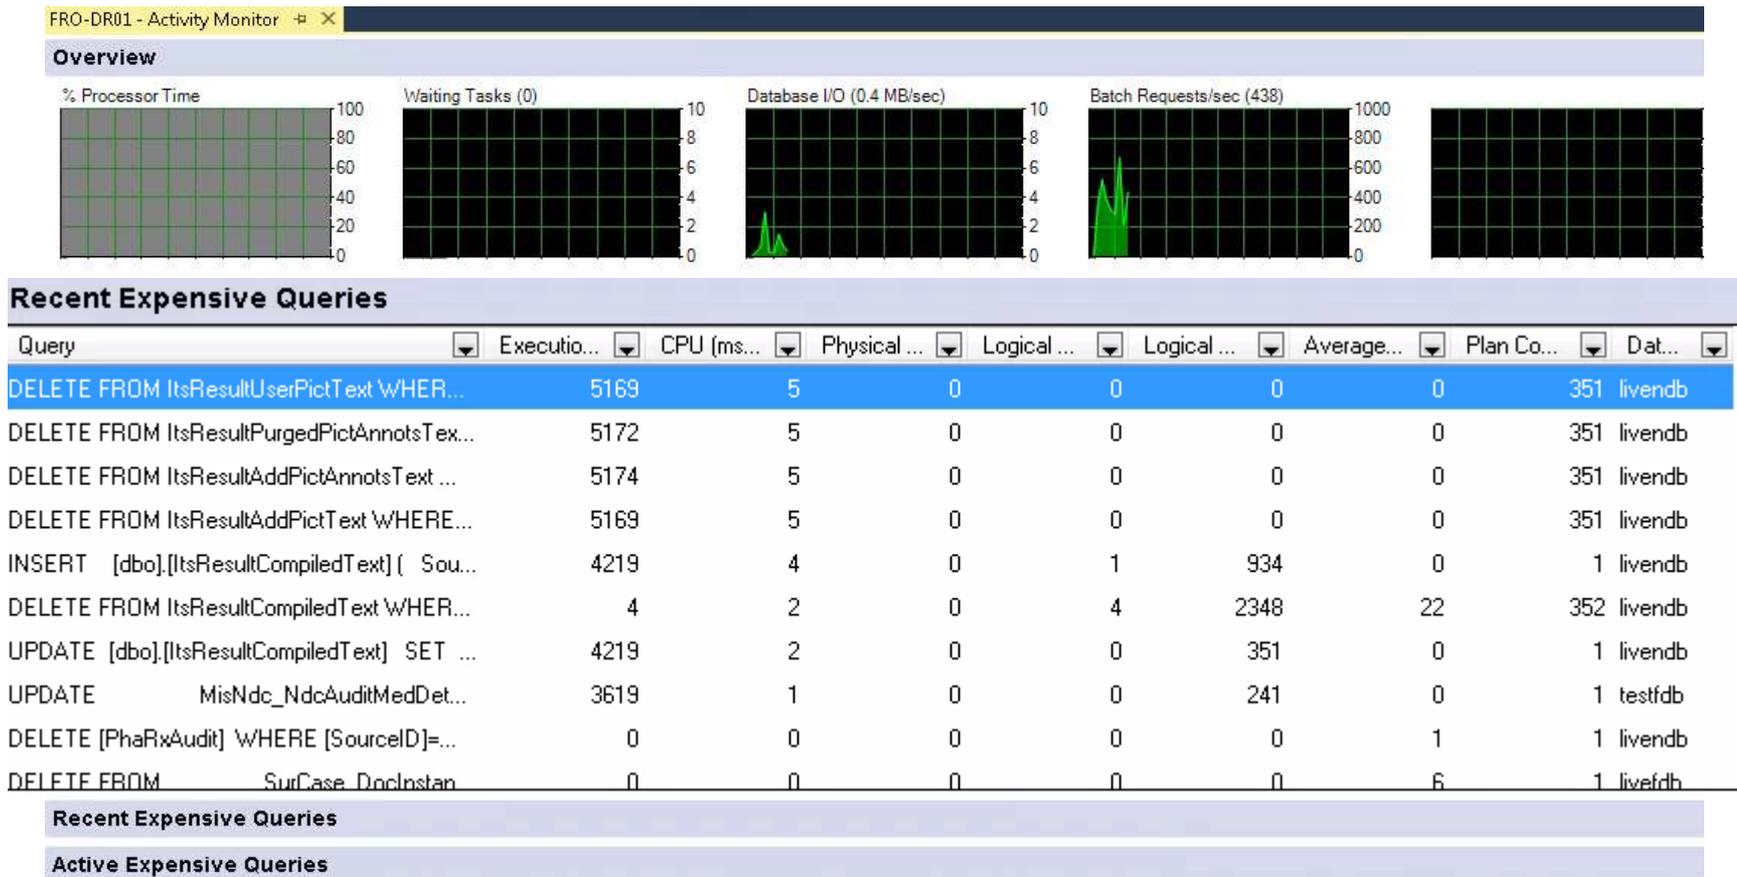
- ALTER INDEX with REORGANIZE for indexes with 5-30% fragmentation
- ALTER INDEX with REBUILD for indexes with >30% fragmentation

<https://docs.microsoft.com/en-us/sql/relational-databases/indexes/reorganize-and-rebuild-indexes?view=sql-server-2016>



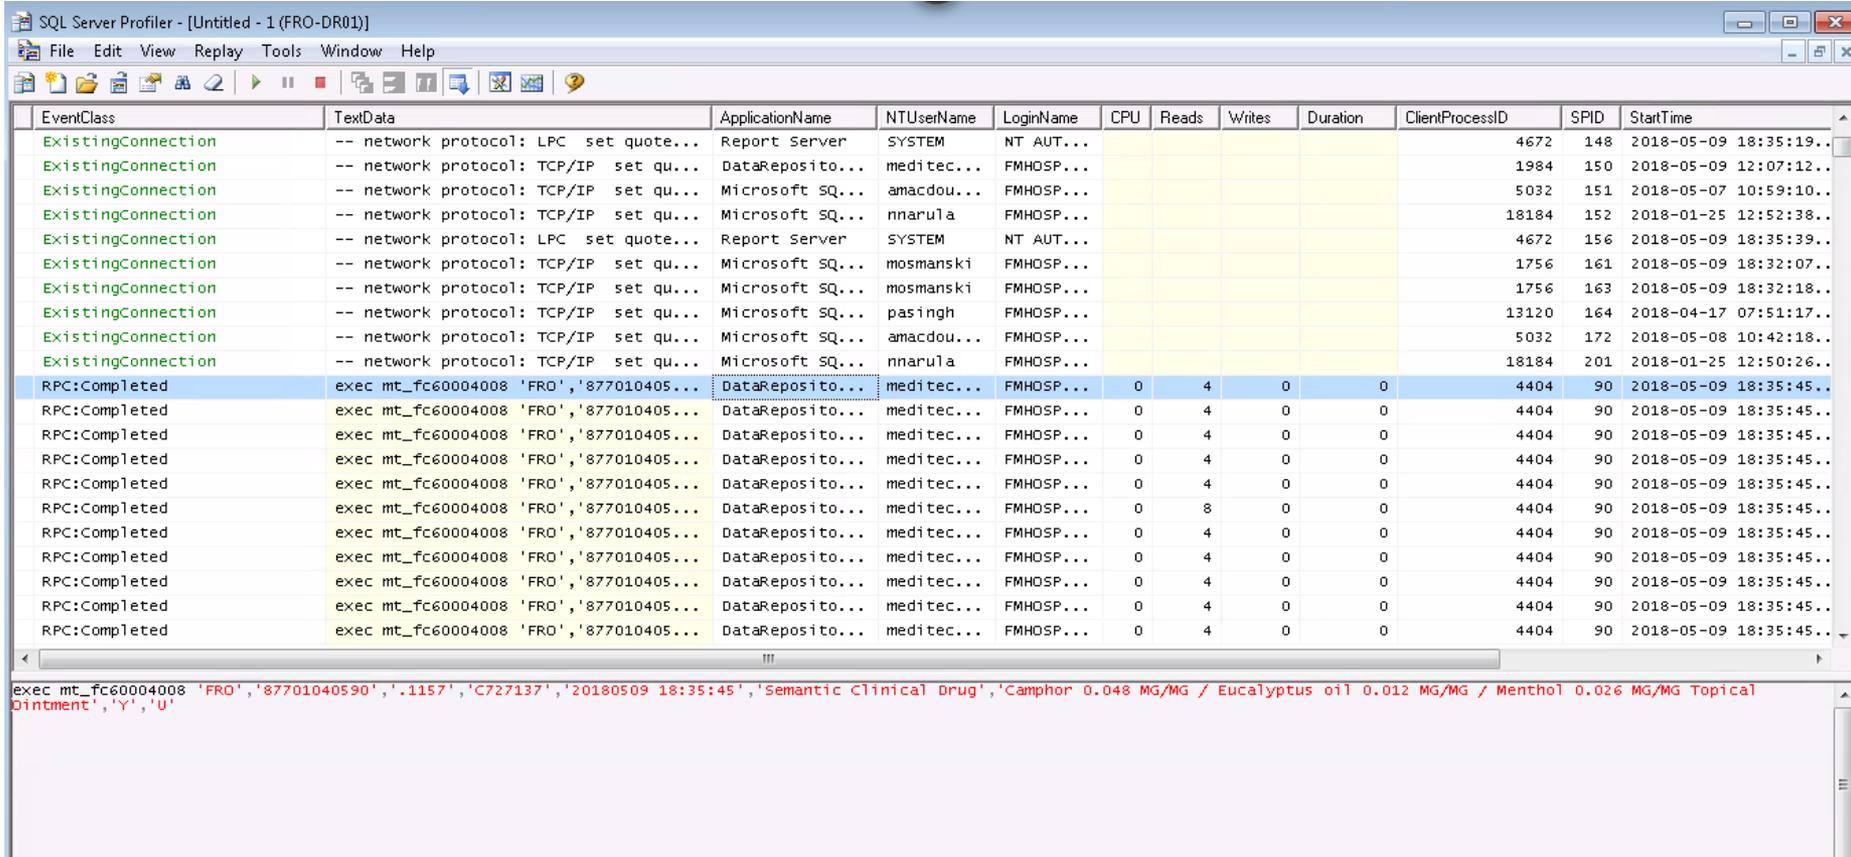
Other performance & monitoring tools

Activity Monitor



Other performance & monitoring tools

SQL Profiler



The screenshot shows the SQL Server Profiler interface with a table of events. The table has columns for EventClass, TextData, ApplicationName, NTUserName, LoginName, CPU, Reads, Writes, Duration, ClientProcessID, SPID, and StartTime. The events include 'ExistingConnection' and 'RPC:Completed'. The 'RPC:Completed' events show a query execution for a medication list.

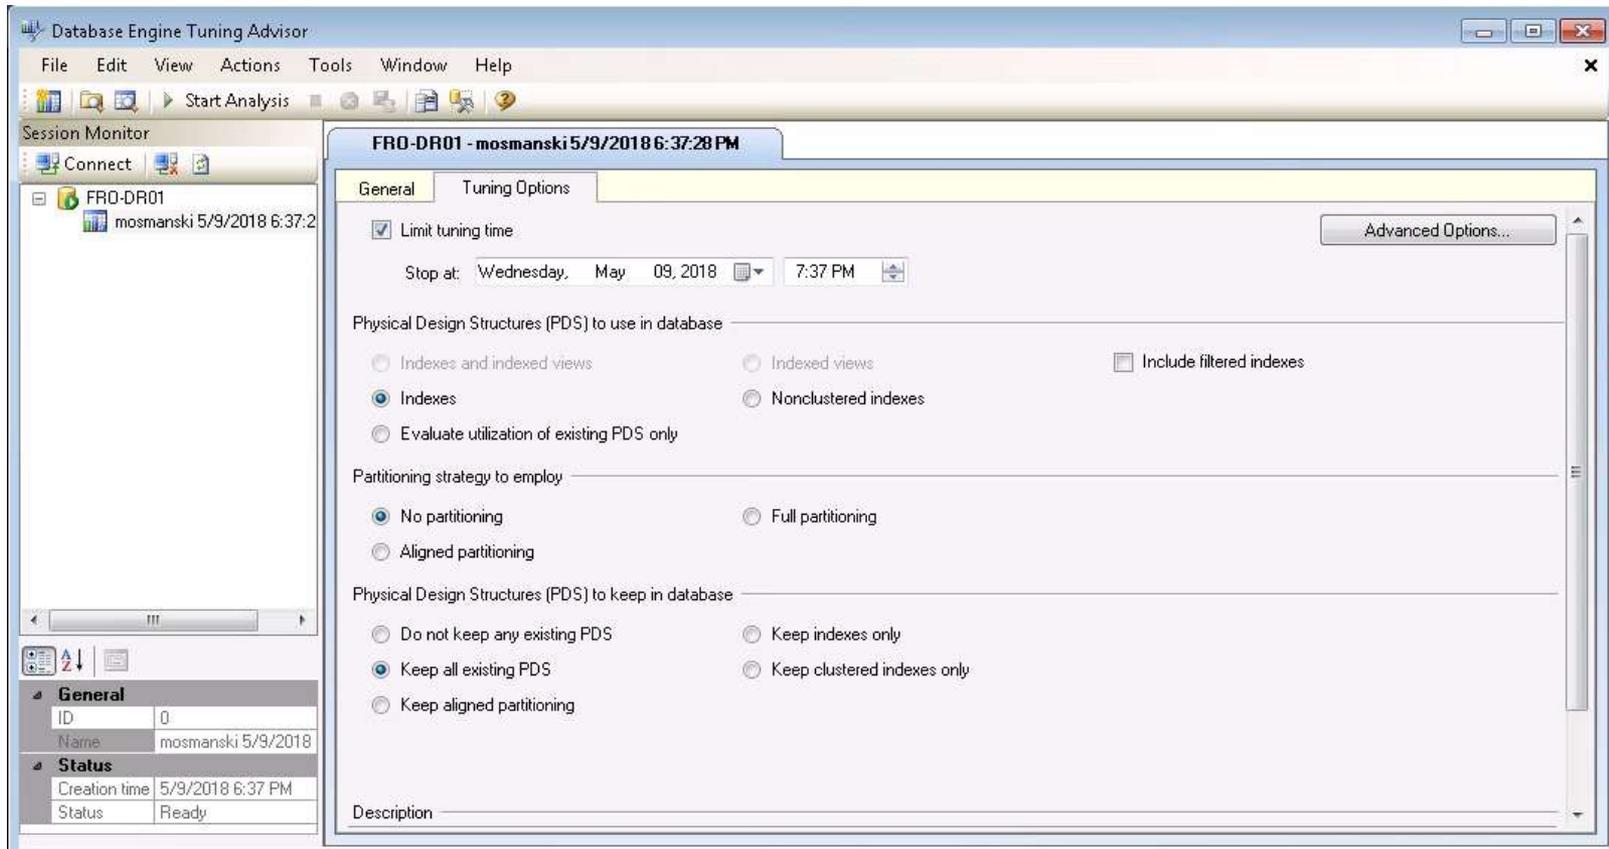
EventClass	TextData	ApplicationName	NTUserName	LoginName	CPU	Reads	Writes	Duration	ClientProcessID	SPID	StartTime
ExistingConnection	-- network protocol: LPC set quote...	Report Server	SYSTEM	NT AUT...					4672	148	2018-05-09 18:35:19..
ExistingConnection	-- network protocol: TCP/IP set qu...	DataReposito...	meditec...	FMHOSP...					1984	150	2018-05-09 12:07:12..
ExistingConnection	-- network protocol: TCP/IP set qu...	Microsoft SQ...	amacdou...	FMHOSP...					5032	151	2018-05-07 10:59:10..
ExistingConnection	-- network protocol: TCP/IP set qu...	Microsoft SQ...	nнарula	FMHOSP...					18184	152	2018-01-25 12:52:38..
ExistingConnection	-- network protocol: LPC set quote...	Report Server	SYSTEM	NT AUT...					4672	156	2018-05-09 18:35:39..
ExistingConnection	-- network protocol: TCP/IP set qu...	Microsoft SQ...	mosmanski	FMHOSP...					1756	161	2018-05-09 18:32:07..
ExistingConnection	-- network protocol: TCP/IP set qu...	Microsoft SQ...	mosmanski	FMHOSP...					1756	163	2018-05-09 18:32:18..
ExistingConnection	-- network protocol: TCP/IP set qu...	Microsoft SQ...	pasingh	FMHOSP...					13120	164	2018-04-17 07:51:17..
ExistingConnection	-- network protocol: TCP/IP set qu...	Microsoft SQ...	amacdou...	FMHOSP...					5032	172	2018-05-08 10:42:18..
ExistingConnection	-- network protocol: TCP/IP set qu...	Microsoft SQ...	nнарula	FMHOSP...					18184	201	2018-01-25 12:50:26..
RPC:Completed	exec mt_fc60004008 'FRO','877010405...	DataReposito...	meditec...	FMHOSP...	0	4	0	0	4404	90	2018-05-09 18:35:45..
RPC:Completed	exec mt_fc60004008 'FRO','877010405...	DataReposito...	meditec...	FMHOSP...	0	4	0	0	4404	90	2018-05-09 18:35:45..
RPC:Completed	exec mt_fc60004008 'FRO','877010405...	DataReposito...	meditec...	FMHOSP...	0	4	0	0	4404	90	2018-05-09 18:35:45..
RPC:Completed	exec mt_fc60004008 'FRO','877010405...	DataReposito...	meditec...	FMHOSP...	0	4	0	0	4404	90	2018-05-09 18:35:45..
RPC:Completed	exec mt_fc60004008 'FRO','877010405...	DataReposito...	meditec...	FMHOSP...	0	4	0	0	4404	90	2018-05-09 18:35:45..
RPC:Completed	exec mt_fc60004008 'FRO','877010405...	DataReposito...	meditec...	FMHOSP...	0	8	0	0	4404	90	2018-05-09 18:35:45..
RPC:Completed	exec mt_fc60004008 'FRO','877010405...	DataReposito...	meditec...	FMHOSP...	0	4	0	0	4404	90	2018-05-09 18:35:45..
RPC:Completed	exec mt_fc60004008 'FRO','877010405...	DataReposito...	meditec...	FMHOSP...	0	4	0	0	4404	90	2018-05-09 18:35:45..
RPC:Completed	exec mt_fc60004008 'FRO','877010405...	DataReposito...	meditec...	FMHOSP...	0	4	0	0	4404	90	2018-05-09 18:35:45..
RPC:Completed	exec mt_fc60004008 'FRO','877010405...	DataReposito...	meditec...	FMHOSP...	0	4	0	0	4404	90	2018-05-09 18:35:45..
RPC:Completed	exec mt_fc60004008 'FRO','877010405...	DataReposito...	meditec...	FMHOSP...	0	4	0	0	4404	90	2018-05-09 18:35:45..

exec mt_fc60004008 'FRO','87701040590','.1157','C727137','20180509 18:35:45','Semantic Clinical Drug','Camphor 0.048 MG/MG / Eucalyptus oil 0.012 MG/MG / Menthol 0.026 MG/MG Topical Ointment','Y','U'



Other performance & monitoring tools

Database Engine Tuning Advisor



Discussion, Q&A



Thank you!