

DAMMIT JIM! DR MCCOY'S FIELD GUIDE TO SYSTEM_HEALTH (AND THE DEFAULT TRACE)

Ed Leighton-Dick



ABOUT ME

- Email: ed@leightondick.com
- Blog: edleightondick.com
- Twitter: [@eleightondick](https://twitter.com/eleightondick)











TABLE OF CONTENTS


1. The Default Trace
2. system_health
3. Ideas
4. Epilogue



CHAPTER 1: THE DEFAULT TRACE


WHAT IS THE DEFAULT TRACE?

- First introduced in SQL Server 2005
- Server-side trace
- Event classes
 - Database
 - Errors and Warnings
 - Full-Text
 - Objects
 - Security Audit
 - Server
- Can be disabled (but not deleted)



MINING THE DEFAULT TRACE

- SQL Server Profiler
- T-SQL
 - `sys.fn_trace_gettable`
 - `sys.trace_events`





SO LET'S JUST KEEP USING THIS...

- Relatively high overhead
- Low-level of detail per event
- Limited number of events
- Based on deprecated method (trace)

- It is still useful... *for now*


CHAPTER 2: SYSTEM_HEALTH

WHAT IS SYSTEM_HEALTH?

- First introduced in SQL 2008
- Uses Extended Events
- Lower resource utilization
- Much more comprehensive
- Events captured in ring buffer
- Lost on system shutdown


"A TRUE BLACK BOX"

- Significantly improved in SQL 2012
- More events
- Subset of events persisted to .xel files
- AlwaysOn_health added




ALWAYSON_HEALTH

- XEvents session for AlwaysOn Availability Groups
- Subset of AG XEvents
- Created with first availability group
- Automatically started on all replicas




WHAT INFORMATION DOES SYSTEM_HEALTH CONTAIN?

- `sp_server_diagnostics`
- Selected errors
- Ring buffer events




SP_SERVER_DIAGNOSTICS

- Special process
 - Internal, preemptive, high-priority thread
 - Reserved memory
 - No locking, synchronization conflicts
- Runs every 5 seconds by default
- Captures data using XEvents
 - System
 - Resource
 - Query Processing
 - I/O
 - Events




ERRORS

- Errors with severity 20+
- Memory-related errors
- Non-yielding scheduler
- Deadlocks
- Latch waits > 15 sec
- Lock waits > 30 sec
- Preemptive waits (external API)
- CLR allocation failures



ERRORS


- Most error events include:
 - Callstack
 - SQL text
 - Session ID
 - More details



RING BUFFER EVENTS


- Memory broker
- Scheduler monitor
- Memory node OOM
- Security
- Connectivity

I haven't found documentation on what any of these contain...



THIS SOUNDS A LOT LIKE THE DEFAULT TRACE...


- Lower resource usage
- Many more events
- Much more detail per event
- Maintained short-term in ring buffer






MINING SYSTEM_HEALTH

- Data stored as binary XML




MINING SYSTEM_HEALTH

- Data stored as binary XML
- ... in the ring buffer
 - 4MB
 - Rolls over frequently




MINING SYSTEM_HEALTH

- Data stored as binary XML
- ... in the ring buffer
 - 4MB
 - Rolls over frequently
- Extract using XQuery




MINING SYSTEM_HEALTH

- Data stored as binary XML
- ... in the ring buffer
 - 4MB
 - Rolls over frequently
- Extract using XQuery
- Minimal documentation




MINING SYSTEM_HEALTH

- Extended Events viewer
 - Not used for ring_buffer target
- T-SQL: sys.dm_xe_*
 - sys.dm_xe_sessions
 - sys.dm_xe_session_targets
 - sys.dm_xe_session_events



UMMM... IT'S NOT THERE...

- system_health can be recreated
- Execute u_tables.sql in the install directory
 - C:\Program Files\Microsoft SQL Server\MSSQLxx.<instanceid>\MSSQL\Install
- Start session after recreating





CHAPTER 3: IDEAS


CHAPTER 4: EPILOGUE

REVIEW

- Default trace: Good, but outdated
- system_health: Lower overhead, more complete
- Both are great troubleshooting tools




RESOURCES



QUESTIONS?

- Email: ed@leightondick.com
- Blog: edleightondick.com
- Twitter: [@eightondick](https://twitter.com/eightondick)



THANK YOU FOR ATTENDING!

~ Please remember to fill out your comment cards ~