

Locks, latches and spinlocks

David Barbarin
SQL Server MVP / MCM




Merci à nos sponsors



Merci à nos volontaires



David BARBARIN

 Microsoft®
SQL Server® ~ depuis 2002
6.5 <= SQL Server <= 2014

 mikedavem1@hotmail.com
david.barbarin@dbi-services.com

 <http://blog.developpez.com/mikedavem> (French)

 <http://tinyurl.com/h5ufofp> (English)

 @mikedavem



Microsoft®
CERTIFIED
Master

dbi services

SQLSaturday Paris 2016

GUSS

A logo for 'PASS SQL Saturday'. It features a green speech bubble icon with a white outline, containing the word 'PASS' in white. To the right of the speech bubble, the words 'SQL' and 'saturday' are written in a bold, green, sans-serif font.



Agenda

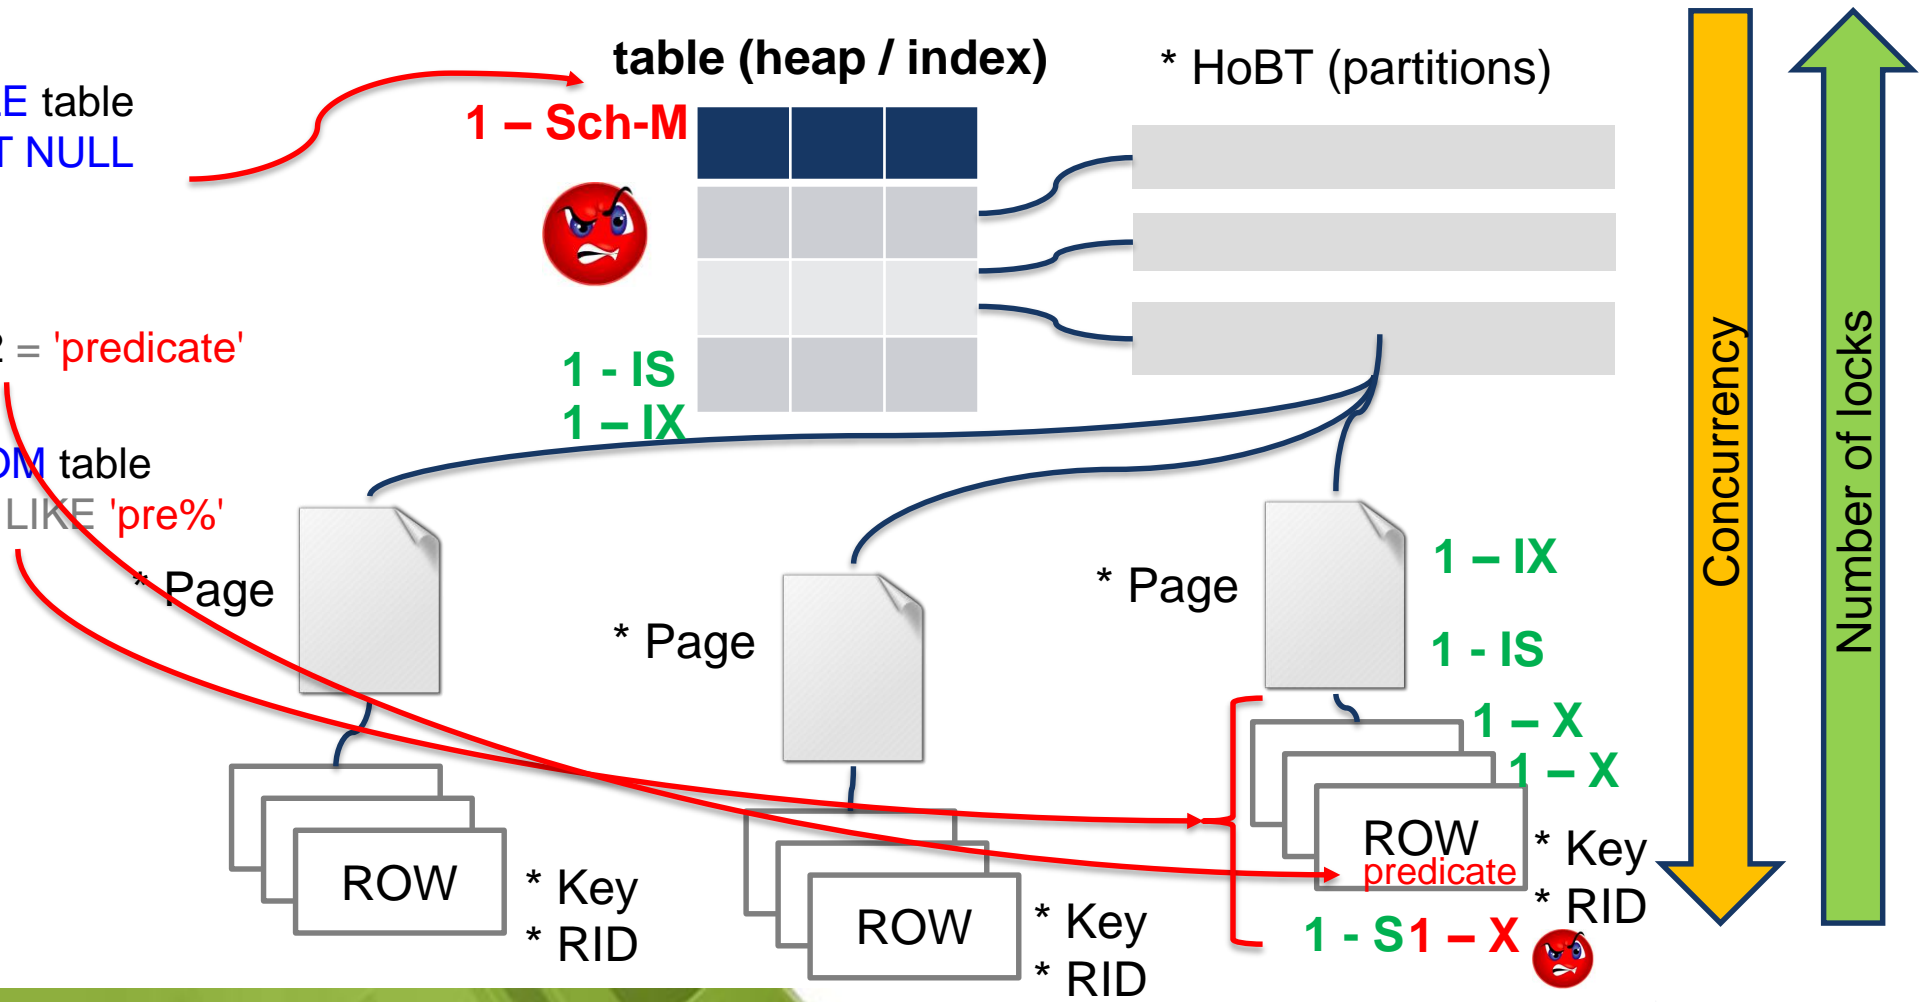


Locks hierarchy

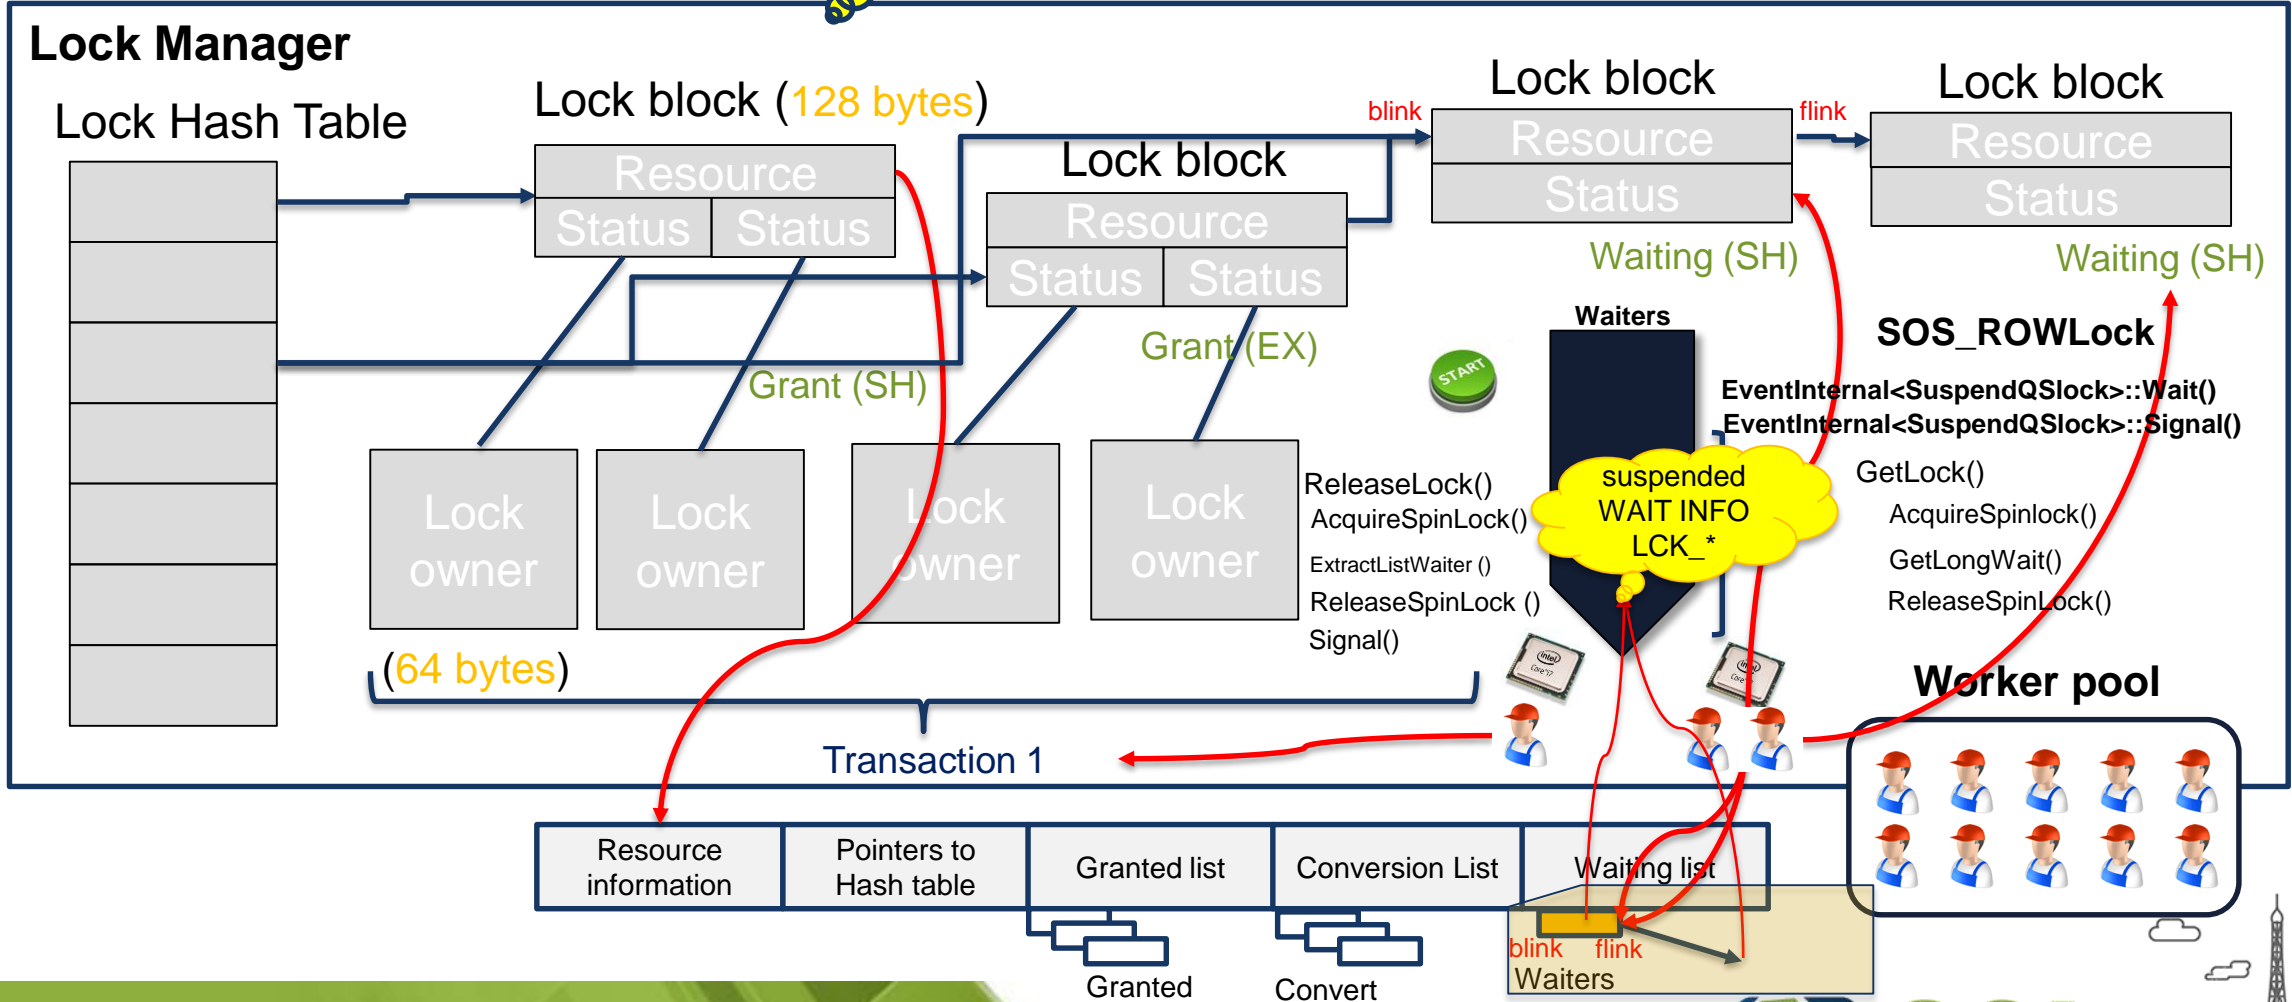
ALTER TABLE table
ADD col3 INT NULL

SELECT col
FROM table
WHERE col2 = 'predicate'

DELETE FROM table
WHERE col2 LIKE 'pre%'



Locks internal



Locks - Demo

Deadlocked schedulers issue

Command - Dump I:\sauv\SQL SERVER\dumps\03 - run out of worker threads cause by tempdb.mdmp - WinDbg:6.3.9600.16384 AMD64

Microsoft (R) Windows Debugger Version 6.3.9600.16384 AMD64
Copyright (c) Microsoft Corporation. All rights reserved.

Loading Dump File [I:\sauv\SQL SERVER\dumps\03 - run out of worker threads cause by tempdb.mdmp]
Comment: 'Stack Trace'
Comment: 'Deadlocked Schedulers'
User Mini Dump File: Only registers, stack and portions of memory are available

***** Symbol Path validation summary *****

Response	Time (ms)	Location
Deferred		srv*c:\symbols*http://msdl.microsoft.com/download/symbols

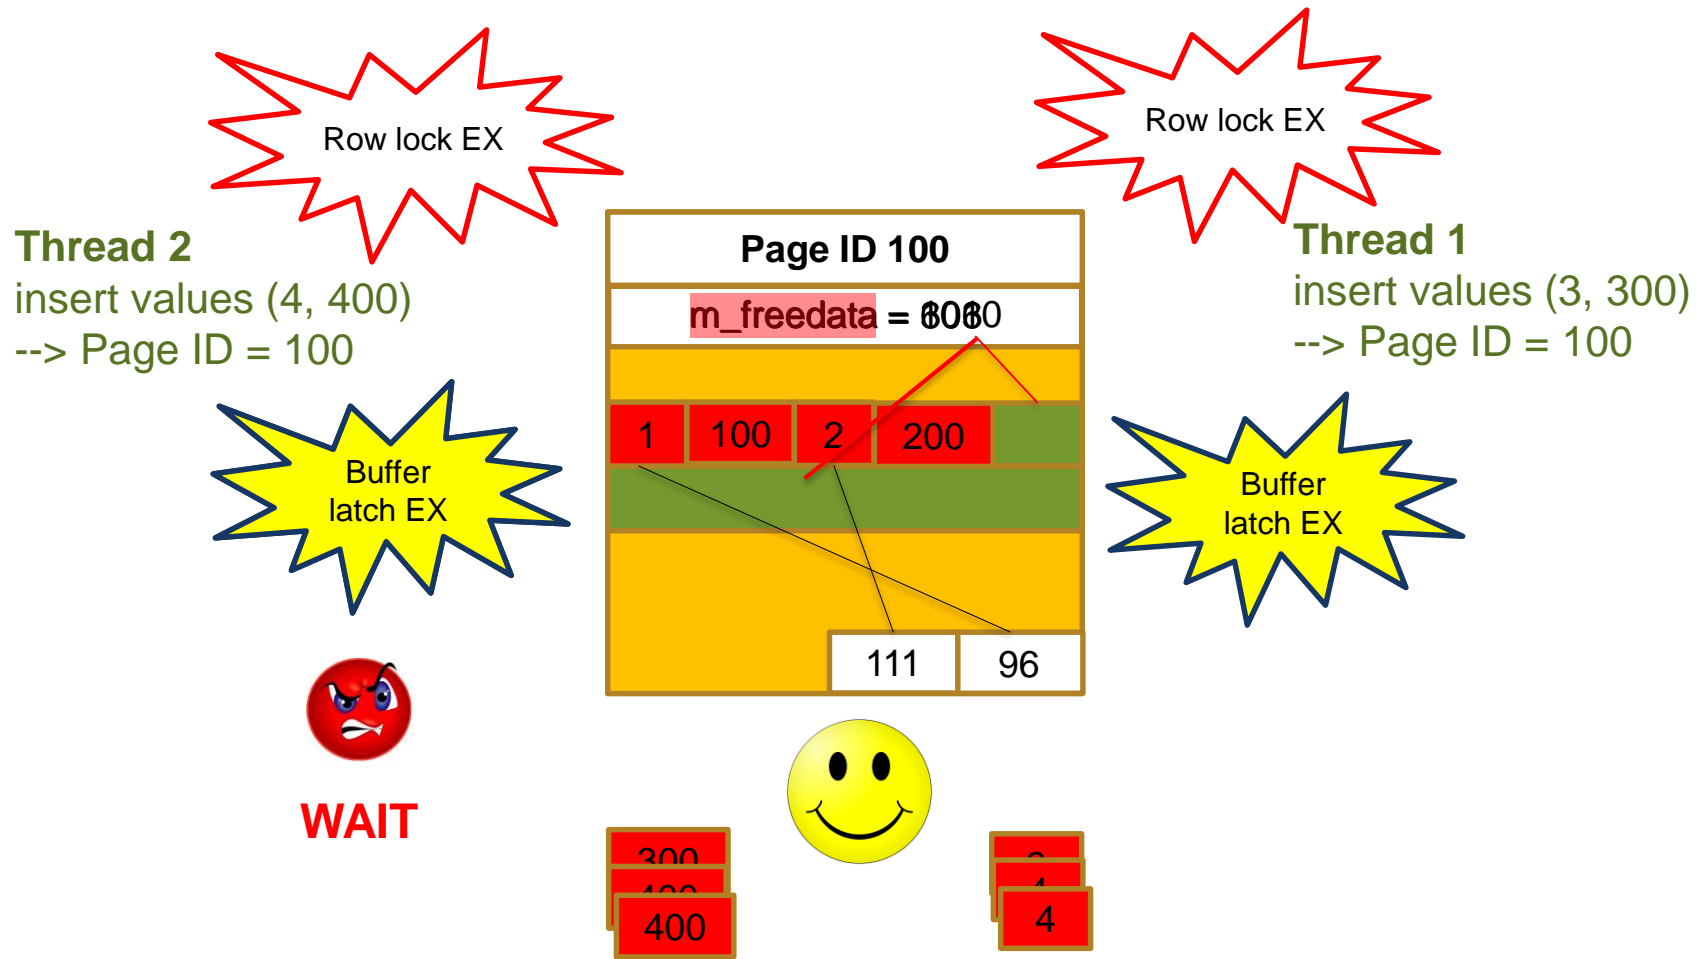
Symbol search path is: srv*c:\symbols*http://msdl.microsoft.com/download/symbols
Executable search path is:
Windows 7 Version 7601 (Service Pack 1) MP (4 procs) Free x64
Product: Server, suite: Enterprise TerminalServer SingleUserTS
Machine Name:
Debug session time: Fri Mar 27 02:05:39.000 2015 (UTC + 2:00)
System Uptime: 53 days 5:59:24.486
Process Uptime: 0 days 0:14:42.000

.....
Loading unloaded module list
.....

This dump file has an exception of interest stored in it.
The stored exception information can be accessed via .ecxr.
(1200.18f8): Unknown exception - code 00000000 (first/second chance not available)
ntdll!NtWaitForSingleObject+0xa:
00000000`77cc12fa c3 ret



Latches – Why are they so important?



Latches

- C++ class to protect important data structure
- Lightweight thread synchronization mechanism
- No isolation level or transaction semantics
- Not directly under user control
- Ownership and waiters
- Compatibility latch matrix
- FIFO based




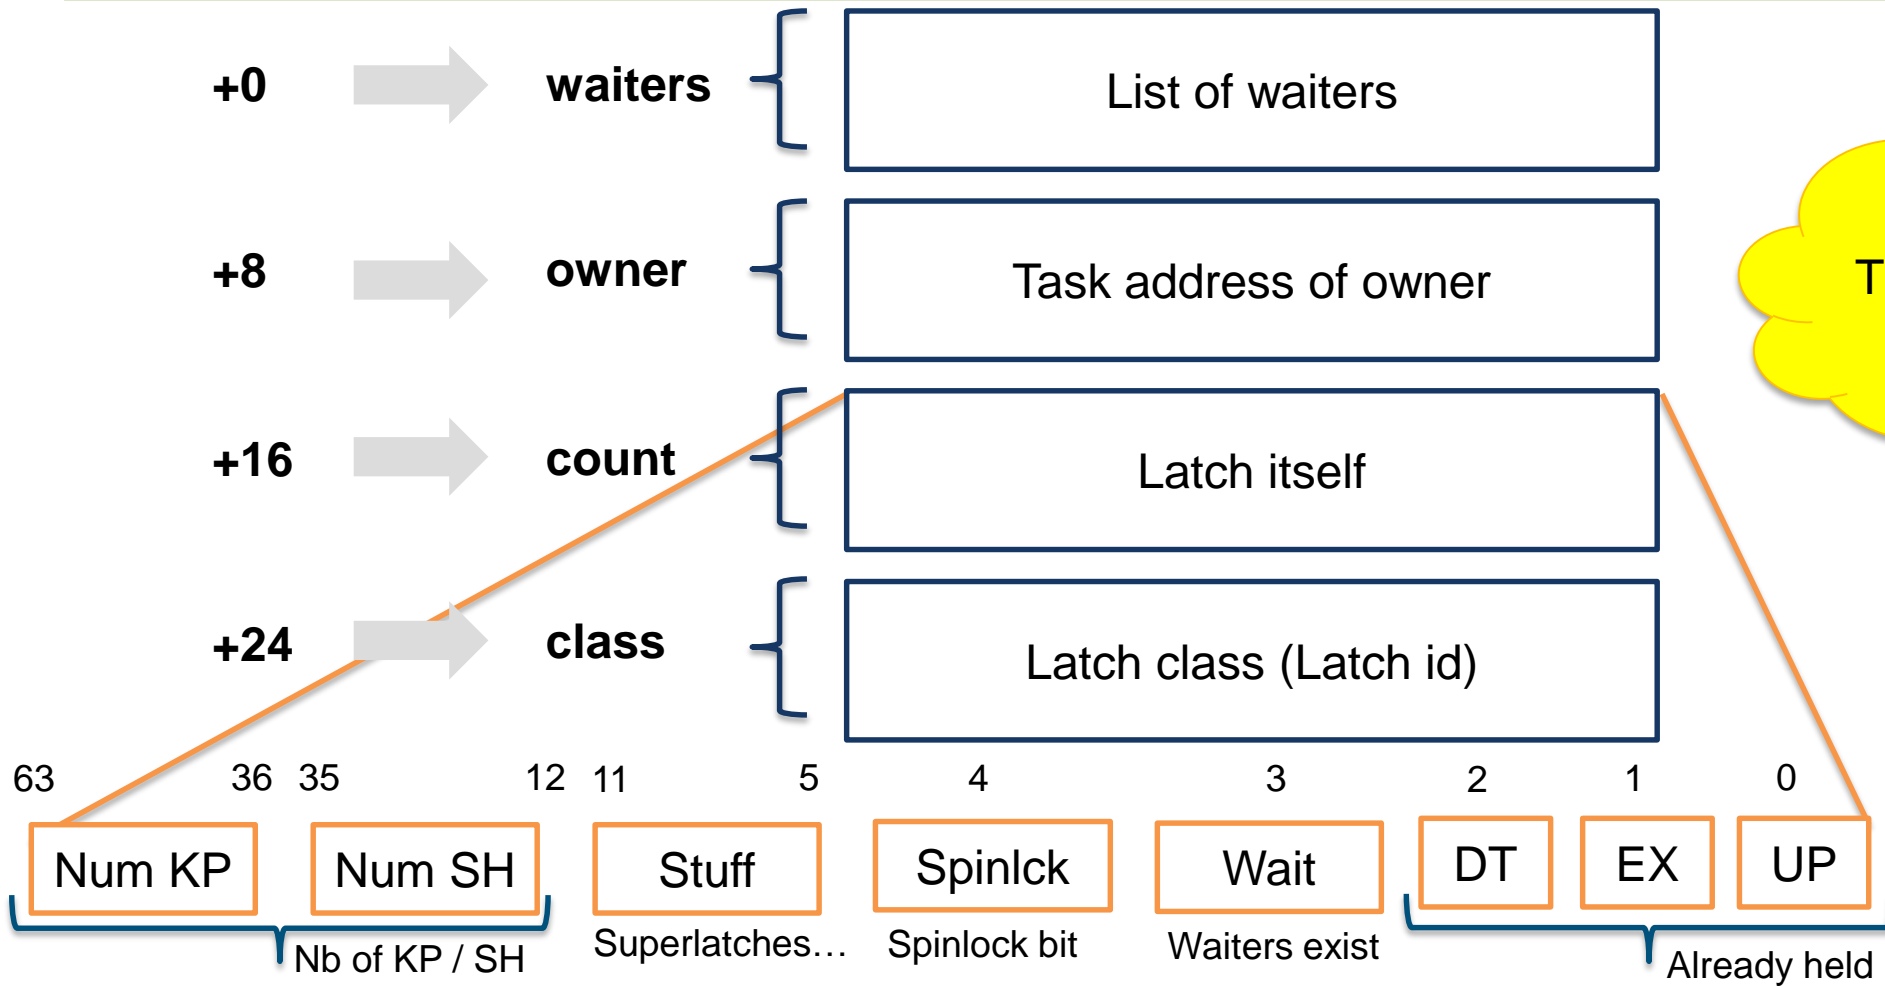
Latches – compatibility matrix

Requested Latch	Currently acquired					
	KP	SH	UP	EX	DT	
KP	Y	Y	Y	Y	N	
SH	Y	Y	Y	N	N	
UP	Y	Y	N	N	N	
EX	Y	N	N	N	N	
DT	N	N	N	N	N	

- KP = Just keeping
- SH = Just reading
- UP = Readers / no writers
- EX = Exclusive / tracking mode
- DT = Destroy (heaviest for freeing)



Latches – C++ structure



Latches - Categories

PAGELATCH_XX
or
PAGEIOLATCH_X
X

Buffer latch (BUF) Latches

- > Protect buffer related to the data page in cache
- > A separate buffer structure to protect data page
- > Use the bstat value to determine if it is an IO page

BUF_ONLRU	0x000001
BUF_DIRTY	0x000002
BUF_IO	0x000004
BUF_HASHED	0x000008
BUF_READ	0x000400
BUF_IOERR	0x000800

LATCH_XX

Non BUF Latches

- > Protect other important memory structures
- > Matrix compatibility is also needed

sys.dm_os_wait_stats
sys.dm_os_latch_stats

sys.dm_exec_requests
sys.dm_os_waiting_tasks
Extended events ...



Latches - Demo

Deadlocked schedulers issue

```
Command - Dump I:\sauv\SQL SERVER\dumps\07 - SQL Server latch timeout.mdmp - WinDbg:6.3.9600.16384 AMD64

Loading Dump File [I:\sauv\SQL SERVER\dumps\07 - SQL Server latch timeout.mdmp]
Comment: 'Stack Trace'
Comment: 'Latch timeout'
User Mini Dump File: Only registers, stack and portions of memory are available

***** Symbol Path validation summary *****
Response                Time (ms)    Location
Deferred                0           srv*c:\symbols*http://msdl.microsoft.com/download/symbols
Symbol search path is:  srv*c:\symbols*http://msdl.microsoft.com/download/symbols
Executable search path is:
Windows 8 Version 9200 MP (12 procs) Free x64
Product: LanManNt, suite: SmallBusiness TerminalServer SmallBusinessRestricted SingleUserTS
Machine Name:
Debug session time: Wed Apr 29 00:06:46.000 2015 (UTC + 2:00)
System Uptime: 1 days 6:53:15.533
Process Uptime: 1 days 6:52:07.000
.....
Loading unloaded module list
.....
This dump file has an exception of interest stored in it.
The stored exception information can be accessed via .ecxr.
(83c.2104): Unknown exception - code 000042ac (first/second chance not available)
ntdll!NtWaitForSingleObject+0xa:
00007ffa`4f337b1a c3      ret
```

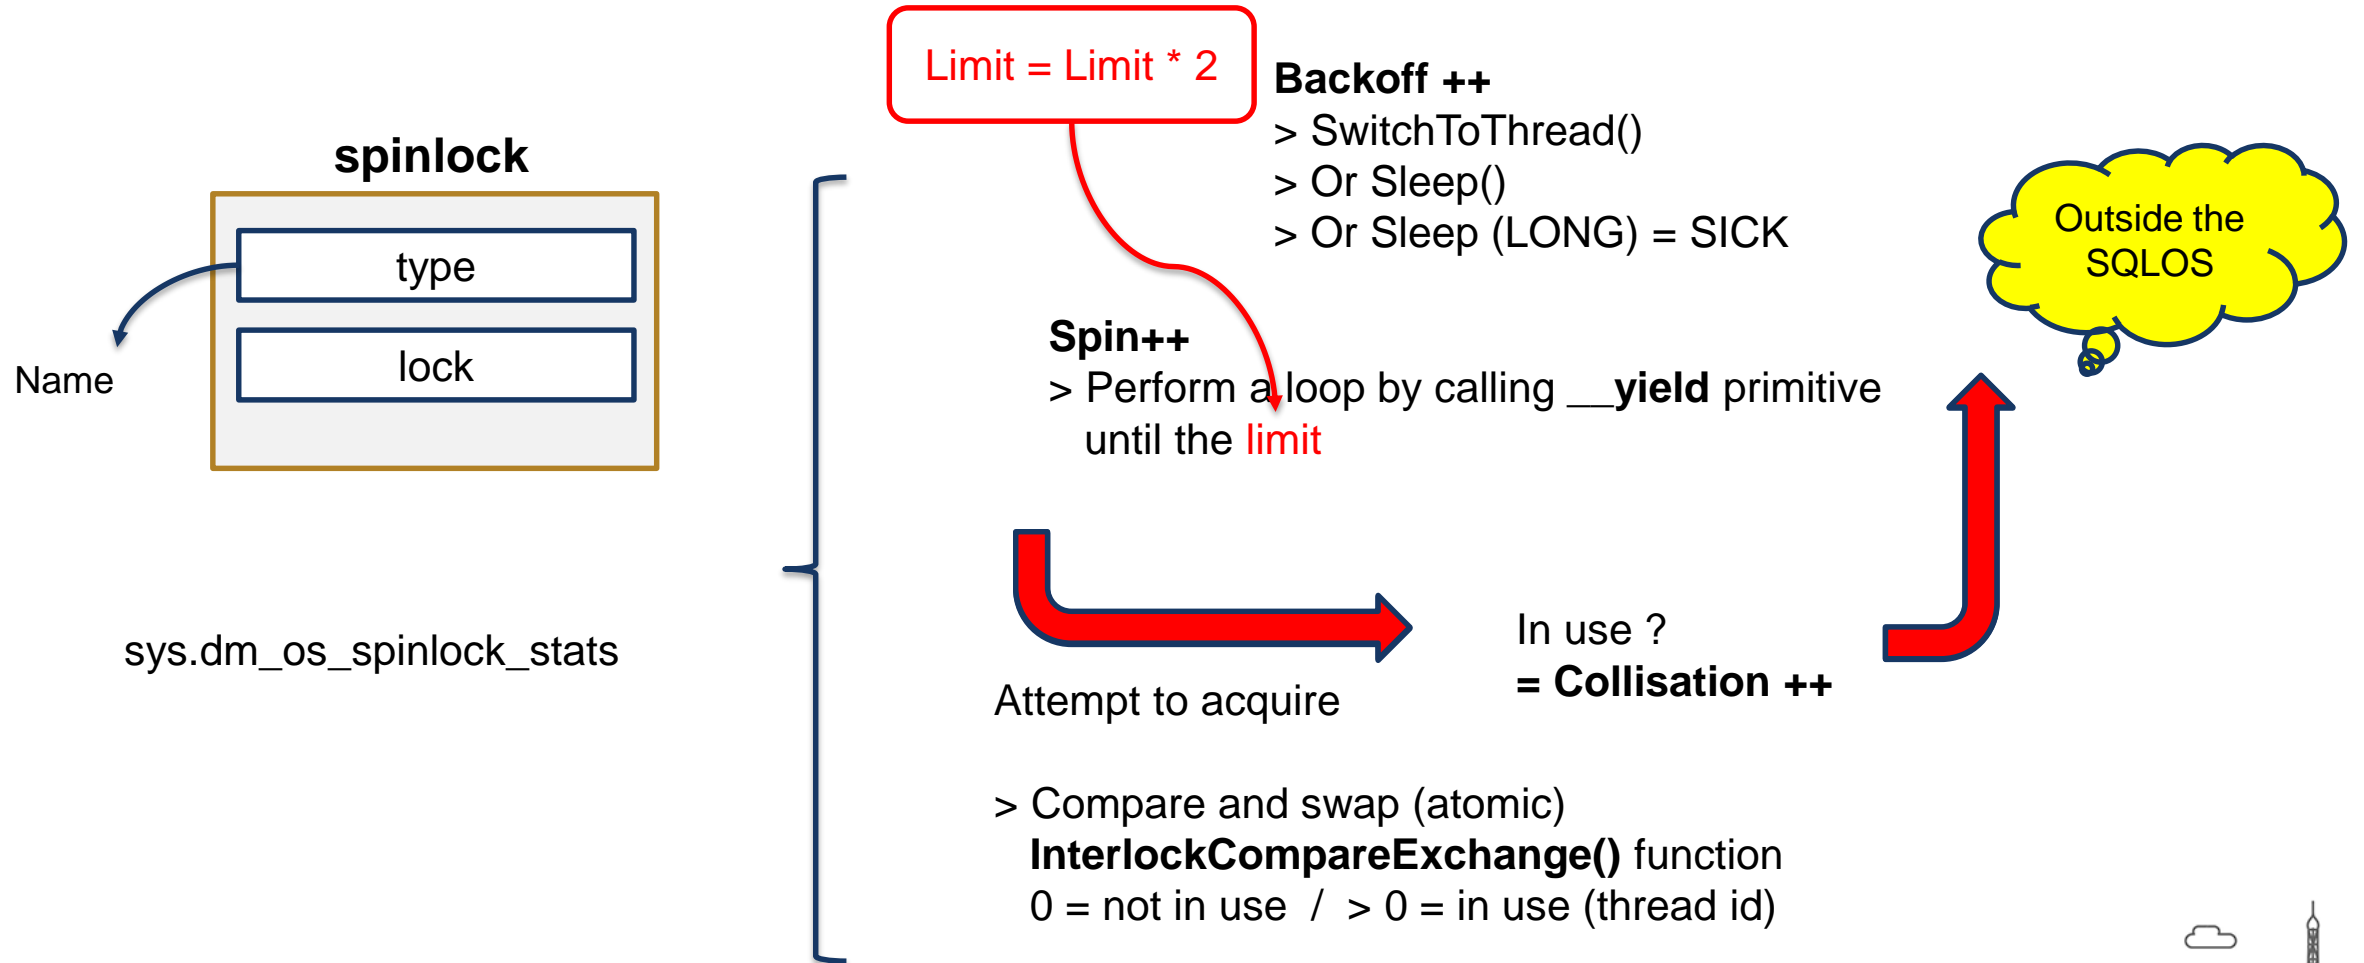


Spinlocks

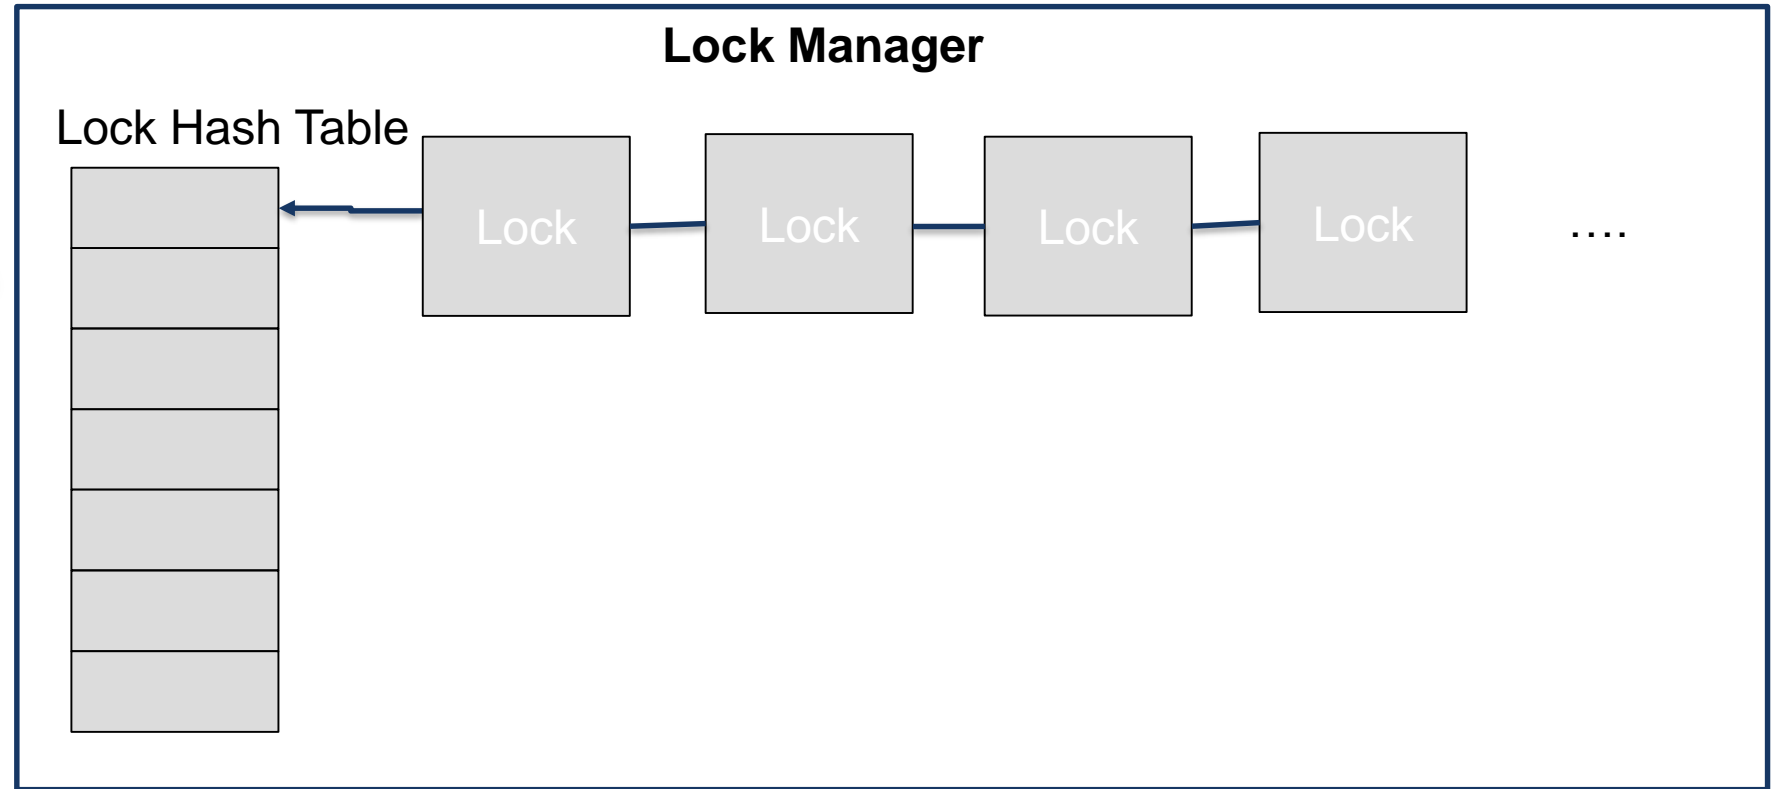
- Latches don't scale very well
- It's a thread safety mechanism in a very lightweight approach
- Mainly used to protect very hot structures lists or hash tables
- Less overhead (Kernel objects are not used)
- No transaction semantic
- No compatibility matrix
- You hold a spinlock or not !



Spinlocks - Principle



Spinlocks - Demo



References

- [MSSQLWIKI](#)
- [Microsoft BOL \(locks\)](#)
- Bob Ward sessions
 - (Inside SQL Server latches / tempdb)
- [Paul Randal blog](#)
- Klaus Aschenbrenner session
 - (Locks, spinlocks and data free data structures)



<http://GUSS.pro/sqlsat>

...Et en plus on peut gagner des cadeaux

